VISUALLY MANAGING IPSEC

THESIS

Peter J. Dell'Accio, 1st Lieutenant, USAF
AFIT/GCO/ENG/10-06

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

AFIT/GCO/ENG/10-06

VISUALLY MANAGING IPSEC

THESIS

Presented to the Faculty

Department of Electrical and Computer Engineering

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the

Degree of Master of Science

Peter J. Dell'Accio, BS
1st Lieutenant, USAF

March 2010

AFIT/GCO/ENG/10-06

VISUALLY MANAGING IPSEC

Peter J. Dell'Accio, B.S. Computer Science
1st Lieutenant, USAF

Approved:

//signed//	8 Feb 10
Lt Col Stuart H. Kurkowski (Chairman)	(8 February 2010)
//signed//	8 Feb 10
Dr. Michael R. Grimaila (Member)	(8 February 2010)
//signed//	8 Feb 10
Timothy H. Lacey (Member)	(8 February 2010)

Abstract

The United States Air Force relies heavily on computer networks to transmit vast amounts of information throughout its organizations and with agencies throughout the Department of Defense. The data take many forms, utilize different protocols, and originate from various platforms and applications. It is not practical to apply security measures specific to individual applications, platforms, and protocols. Internet Protocol Security (IPsec) is a set of protocols designed to secure data traveling over IP networks, including the Internet. By applying security at the network layer of communications, data packets can be secured regardless of what application generated the data or which protocol is used to transport it. However, the complexity of managing IPsec on a production network, particularly using the basic command-line tools available today, is the limiting factor to widespread deployment. This thesis explores several visualizations of IPsec data, evaluates the viability of using visualization to represent and manage IPsec, and proposes an interface for a visual IPsec management application to simplify IPsec management and make this powerful security option more accessible to the information warfighter.

Acknowledgments

I would like to thank my advisor Lt Col Stuart Kurkowski. His time and guidance were essential to my being able to complete this effort.

I would also like to thank the members of my committee, Dr. Michael Grimaila, and Tim Lacey for taking the time to review and critique my efforts.

I was also fortunate to be able to collaborate with a few additional individuals on this effort as well. I wish to thank Jonathan Alley for his help in developing the simulation presented in this thesis. I would also like to thank LTC Greg Conti, author of "Security Data Visualization", and Jon Snader, author of "VPNs Illustrated: Tunnels, VPNs, and IPsec". These individuals provided valuable feedback and insights regarding visualization and IPsec respectively, and I appreciate them taking time to support my efforts.

Lastly, I would like to thank my family for dealing with another chapter in my military career that required the diversion of a considerable amount of my time and attention.

Peter J. Dell'Accio

Table of Contents

List of Figures

Figure Page

List of Tables

Table Page

VISUALLY MANAGING IPSEC

I. Introduction

In today's environment, the information warfighter is as much at the tip of the spear as the troops on the ground or the pilot putting bombs on target. The continuing evolution of an Air Force Cyber Command, currently as a Numbered Air Force, and the establishment of a broader U. S. Cyber Command (USCYBERCOM) illustrate the ever-growing importance of, and focus on, our information and networked communications. As always, the security of our information and communications channels is of paramount importance. The wide range of applications, data types, protocols, standards (or lack of standards in some cases), and technologies makes network security a challenge to say the least.

Internet Protocol Security (IPsec) [1] was designed by the Internet Engineering Task Force (IETF) to secure individual data packets at the network layer of communications by applying cryptographic techniques that enable data origin authentication, connectionless data integrity, data confidentiality, traffic analysis protection, and replay protection. Since IPsec is applied to raw IP packets, the information can be protected regardless of what application generated the data or what transport layer protocol is being used. Additionally, IPsec can be applied using various levels of granularity. A simple rule could be applied to all nodes in a given network requiring all information transmitted within that network be encrypted. In contrast, each

individual node could be assigned a distinct IPsec policy defining which hosts it can communicate with, over which ports, and using which protocols and cryptographic algorithms.

As the number of nodes in a network increases, the complexity of managing IPsec rules and keys increases as well. With visualization, there is the potential to improve situation awareness and make IPsec more manageable by providing an intuitive interface and developing powerful functionality that would not be feasible otherwise. Since IPsec rules essentially represent logical communication paths between network nodes, simplifying IPsec management seems the perfect driver problem for developing a visualization schema that could be easily extended to other similar networking applications.

1.1. Problem Statement

The problem addressed in this thesis is the lack of a way to effectively manage IPsec on production networks. The difficulty of managing complex IPsec rule sets on networks with more than a few nodes seems to be the main factor limiting deployment of this powerful, versatile network security tool.

1.2. Research Goals

The initial goal of this research is to develop techniques, primarily through visualization, that simplify IPsec management. By streamlining the implementation of IPsec and employing existing techniques for visually managing computer networks we can lay a solid foundation for building a straightforward, intuitive IPsec management tool. By developing an approach to visualizing IPsec rules, we can provide

administrators with a way to keep track of and manage complex rule sets on production networks. A second goal in visualizing IPsec rules is to develop a visualization that is immediately familiar and intuitive to network administrators rather than producing a stand-alone visualization that may lack any intuitive context. Finally, we aim to develop a visualization schema that not only addresses the driver problem of simplifying IPsec management but can be applied to other applications for visualizing logical relationships between network nodes.

1.3. Assumptions

This thesis assumes the following:

1. The approaches to visualizing IPsec deployed on a network are independent of the specific IPsec implementation and the network it is running on (hardware platforms, operating systems, protocol versions, etc).
2. There exists a suitable framework for deploying an IPsec management tool to nodes throughout a local area network.
3. IPsec management functionality can be incorporated into an existing network management tool for system discovery and network visualization. The visualizations presented here can be adapted to fit the aesthetics of a specific network management tool regarding icons, color scheme, etc.
4. It is possible to develop a tool that allows for remote IPsec administration similar to existing IPsec management tools and that can use the IPsec information on each system to render the visualizations suggested by this thesis.

1.4. Contributions

This research lays the foundation for an IPsec management tool that leverages the benefits of visualization techniques to make IPsec more accessible and easier to manage beyond deployment on a small network or in a laboratory environment. It clearly quantifies the potential impact of streamlining the IPsec protocol and presents new directions in management functionality beyond what existing IPsec management tools provide. Beyond IPsec, this research presents a novel way of managing relationships between network nodes. Regarding visualization, this research applies existing visualization techniques to a class of problem and evaluates their applicability to that area. Additionally, unique approaches to visualization issues and concepts are presented and explored.

1.5. Thesis Organization

Chapter II provides background information on the current technologies that allow the IPsec management tool proposed by this thesis to be realized. Section 2.1 discusses data visualization, including discussions specifically about visualization techniques, network visualization, and situation awareness. Section 2.2 provides an overview of IPsec, identifying the various protocols, operating modes, cryptographic techniques, and other mechanisms IPsec uses to provide network-layer security.

Chapter III identifies the specific areas this thesis focuses on in attempting to simplify IPsec management. Section 3.1 examines the various parts of IPsec. We will look for ways to limit complexity by reducing the number of options presented to an administrator without limiting the effectiveness of IPsec. Section 3.2 explores

approaches evaluated for visualizing IPsec rules and presents the approach and a suggested interface as the foundation for developing an IPsec management tool.

In Chapter IV we present the results of the design decisions and development efforts laid out in Chapter III. In Section 4.1 we identify the impact of the decisions regarding streamlining IPsec on both administration and visualization. Section 4.2 provides several potential views based on the schema for visualizing IPsec rules identified in 3.2. This allows us to explore the viability of the proposed visualization with the specific goal of simplifying management of complex IPsec rule sets in mind. This section also describes some of the potential functionality that became apparent as the visualization evolved. Section 4.3 addresses several issues that came to light as this research effort progressed.

Chapter V summarizes the findings of this thesis and identifies how intended research goals were met and the contributions this research provided. Finally, we suggest potential areas for future research in developing and refining the tool presented by this thesis.

II. Background Information

In this chapter, we discuss the concepts behind visualization that make it ideal for addressing complex problems involving large amounts of data. We will explore the benefits of visualization to network management, look at several different visualization techniques, and examine the impact and importance of visualization on situation awareness. We will then explore the fundamental building blocks of IPsec and how they provide network-layer security for packets flowing across IP networks, including the Internet. This information will provide the foundation upon which a more usable, yet robust, IPsec management tool can be built.

2.1. Visualization

A picture is worth a thousand words. The simple truth behind this cliché has driven a substantial amount of study and research into various areas of visualization. Human beings are visual creatures. Though we certainly collect information through all of our senses, the human visual system takes in more data than the other senses combined [2].

Text is a simple form of visualization in itself. The printed number "5" is nothing more than a simple way to visualize a specific quantity of something. However, it becomes difficult to convey information quickly as the amount of text increases. For example, consider an attempt to determine how the Democrat versus Republican popular vote of the 2008 U.S. presidential election played out. Given a table that simply listed the states and which party won the majority vote in each state, it may be feasible to navigate through the data and find the information of interest. That is, since the table

would be relatively small, any information searched for, such as the total number of states for each party or how particular states voted, would likely be found fairly quickly. However, examine Figure 1, which presents the same information visually. Not only can the same information the table would provide be located quickly, but additional information can be inferred by groupings of colored regions.

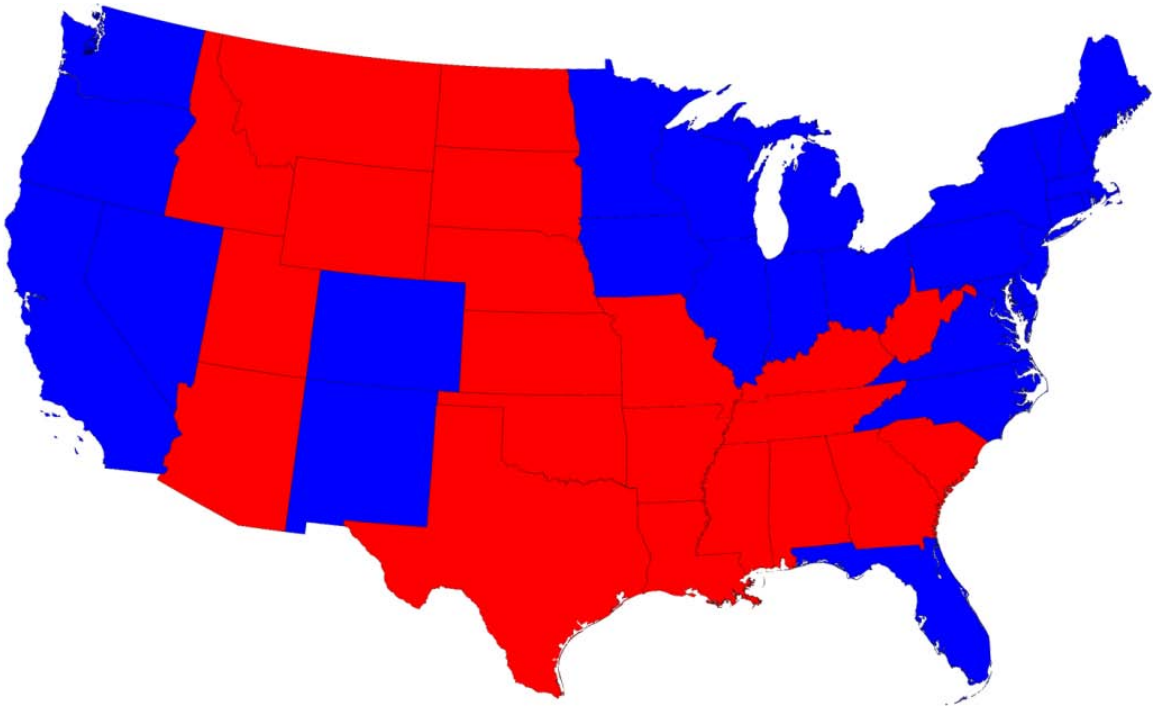

Figure 1: Breakdown by state of 2008 U.S. presidential election popular vote (red = Republican, blue = Democrat) [3]

Now consider if the information was desired by county rather than by state. Additionally, consider more granularity than simply which party had the majority was required. A table would require 3,115 rows to present the results for each county, and the

percentage of votes each party garnered would need to be included. Figure 2 presents this additional information (a legend would need to be included associating percentages to shades of purple) using no more real estate than in Figure 1.

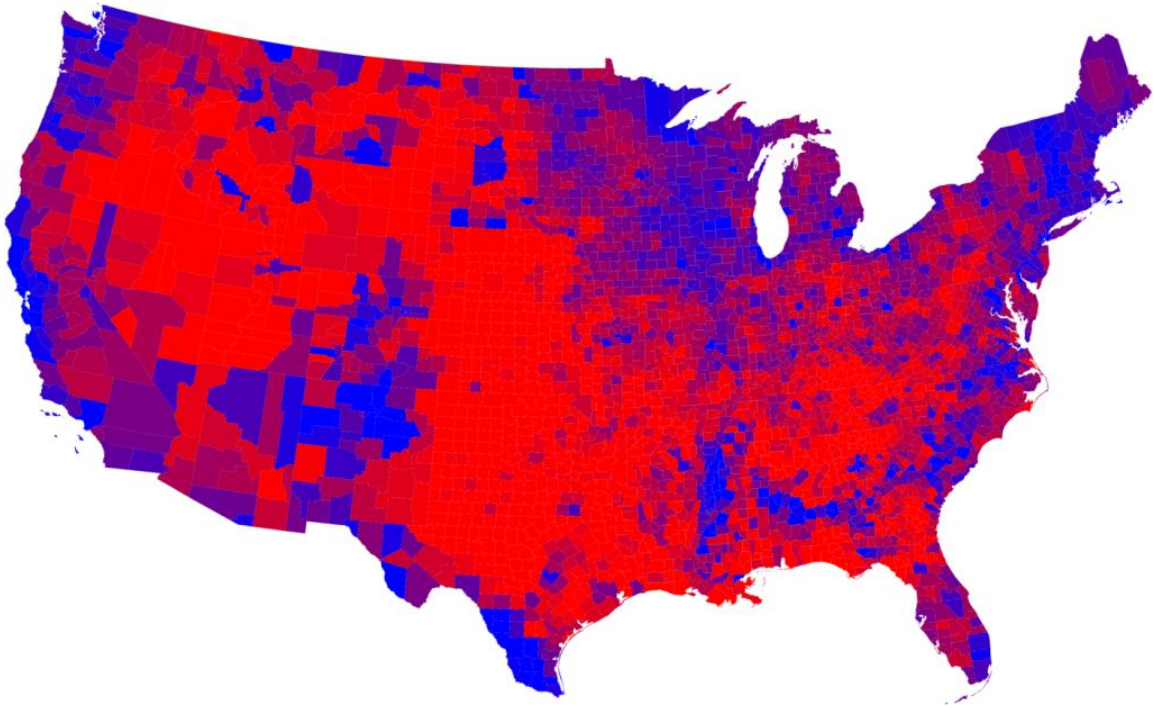

Figure 2: Breakdown by county of 2008 U.S. presidential election popular vote (red = Republican, blue = Democrat) [3]

To be fair, neither approach is perfect. A simple table arranged alphabetically and by state might be better suited for locating the information pertaining to specific counties since the by-county map cannot be labeled practically. However, as alluded to above, information such as the concentration of a given party in a specific area can quickly be discerned from the visualization but would be tedious to try to determine from a table.

Also, when displayed by a computer, the power of the visualization can be enhanced further in ways that would correlate to additional columns of text in a table.

Visualization draws some of its strength from the concept of preattentive processing [2] [4] [5]. Preattentive processing refers to our ability to visually identify and distinguish between certain features without focused attention or conscious effort. There are several features that are preattentively processed and can be exploited to increase the effectiveness of a given visualization. These features can be grouped into four general categories: form (size, shape, grouping, etc.), color (hue, intensity), motion (flicker, direction), and spatial position (2D position, depth, convex/concave shape). Figure 3 illustrates some of the preattentively processed features.

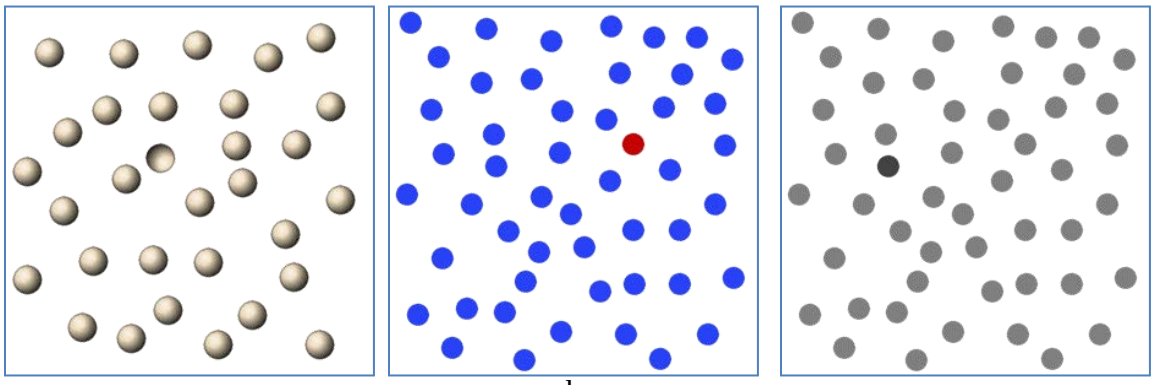

a. b. c.

Figure 3: Examples of preattentively processed features using a.) concavity, b.) color, and c.) intensity [6]

For most, it should be almost automatic to identify the outliers in each example in Figure 3. Techniques such as these are often applied to visualizations to identify anomalies, significant events, like objects, etc. A common example would be showing online systems as green and offline systems as red in network visualizations.

It is important to keep in mind when employing visualization that, like most anything else, too much of a good thing can be bad. A poorly designed visualization can potentially be even more difficult to use than the raw data it attempts to convey. Examine Figure 4, which shows a program called Grokker that presents an alternate view of Yahoo! search results about visualization based on date and Yahoo! ranking [7].

Figure 4: Representation of Yahoo! search results using Grokker [7]

The visualization in Figure 4 includes several preattentive features, including shape, size, color, and grouping. However, these features cannot be used for their preattentive properties because there are simply too many of them. The viewer cannot quickly identify objects based on color due to the number of colors displayed. Similarly, there seems to be a fairly even distribution of circles and squares and a great many sizes, where intuitively, a viewer would expect there to be some significance to the variances. If these features were used for their preattentive properties, the viewer would need to learn and track the different meanings behind the various sizes, shapes, and colors in addition to how results are nested. Simply stated, the visualization is not intuitive. The main subcategories seem to be labeled alphabetically (except for the one labeled 'More…'). The rest of the results are not labeled except for one (the 'Software' container inside the 'Data' container), implying the objects within are results relating to their respective containers. This requires the viewer to either mouse over each entry or drill down into each container to view the results, but there is no implied order or structure to guide the user. An expandable list view organized by main category would likely have been more familiar and simpler to navigate for most viewers.

2.1.1. Visualization Techniques

There are many different ways to visualize information. From simple pie charts and line graphs to today's advanced imaging and display technologies, visualization is limited only by imagination. Finding the right way to visualize a particular set of information can be a challenge. The techniques described below provide an overview of some of the more common approaches used in information visualization.

2.1.1.1 *Charts and Graphs.* Charts and graphs are perhaps the most common and familiar ways to graphically represent data. Pie charts are useful for comparing values that represent a percentage of a whole, bar charts depict a given data dimension as a count, and line charts are suited for displaying data to identify trends [5]. There are many variations of these simple charts including stacked and 3D versions and histograms. Graphs have of course evolved over time to meet different data visualization needs. A good example is the scatter plot, which is common in network security and is used to detect trends and examine relationships between data points [5]. Figure 5 provides examples of some of these simple data visualizations.

Figure 5: Examples of various simple graphs [8]

There are also several techniques geared towards displaying multivariate data such as small multiples [9] [10] and heat maps [10]. Another multivariate display is the parallel coordinate graph [5]. In a parallel coordinate graph, multiple axes are drawn equidistant from each other to represent individual data points. The values for each dimension are plotted along the appropriate axis, and a line spanning the axes represents an individual data entry. The Parallel coordinate graph in Figure 6 shows several data points for over 3,000 U.S. counties at once.

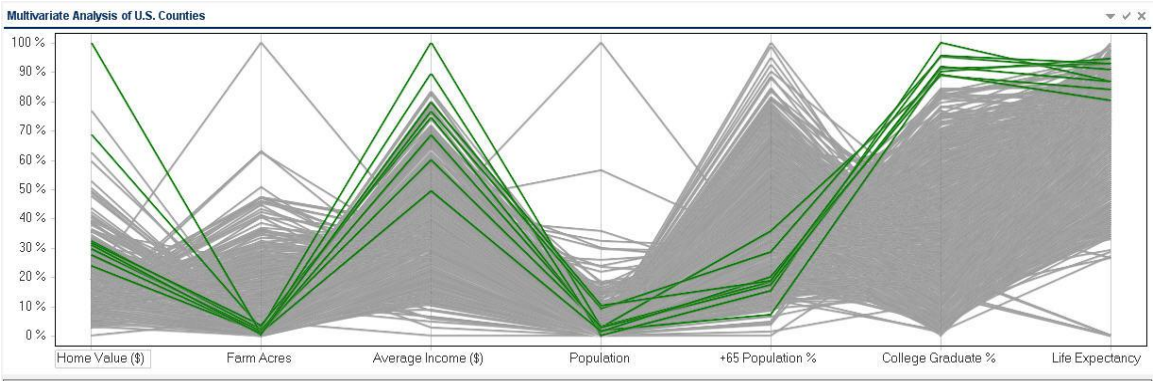

Figure 6: Parallel coordinate graph displaying seven data points of U.S. counties

The sheer number of data entries in such a confined space creates a mostly indistinguishable mass of line plots that would seem to render such a graph useless. However, the intent is not to enable the viewer to immediately know the information pertaining to each individual county. The true value of these graphs lies in the interface that allows a user to select specific points of interest and filter out the rest of the data. The graph in Figure 6 highlights the eight counties with the highest percentage of college

graduates. Parallel coordinate graphs can present a huge amount of information in a small space and quickly identify trends in the data, and filters can allow users to quickly isolate any type of information from massive datasets to analyze specific items of interest.

Link graphs, or link-node graphs, are particularly well-suited for visualizing relationships between network nodes [5]. Physical communications paths, logical connections, and various data dimensions such as data type or bandwidth are some examples of the network information link graphs typically represent. Each physical network node is represented by a node on the graph. Any specified relationship between two nodes is represented by a link or edge. Information can be encoded into both the nodes and the links using size, color, shape, line thickness, or any other distinguishable feature.

Link graphs have been used to show relationships between entities beyond networking and communications. Figure 7 [12] shows a Facebook Friend Wheel, which is a radial link graph displaying all of an individual's friends plotted as nodes on a circle. Any person on the circle who shares a friend in common with that individual is connected to the other person with a link. Color in Figure 7 has no apparent informational value and is used for aesthetics only.

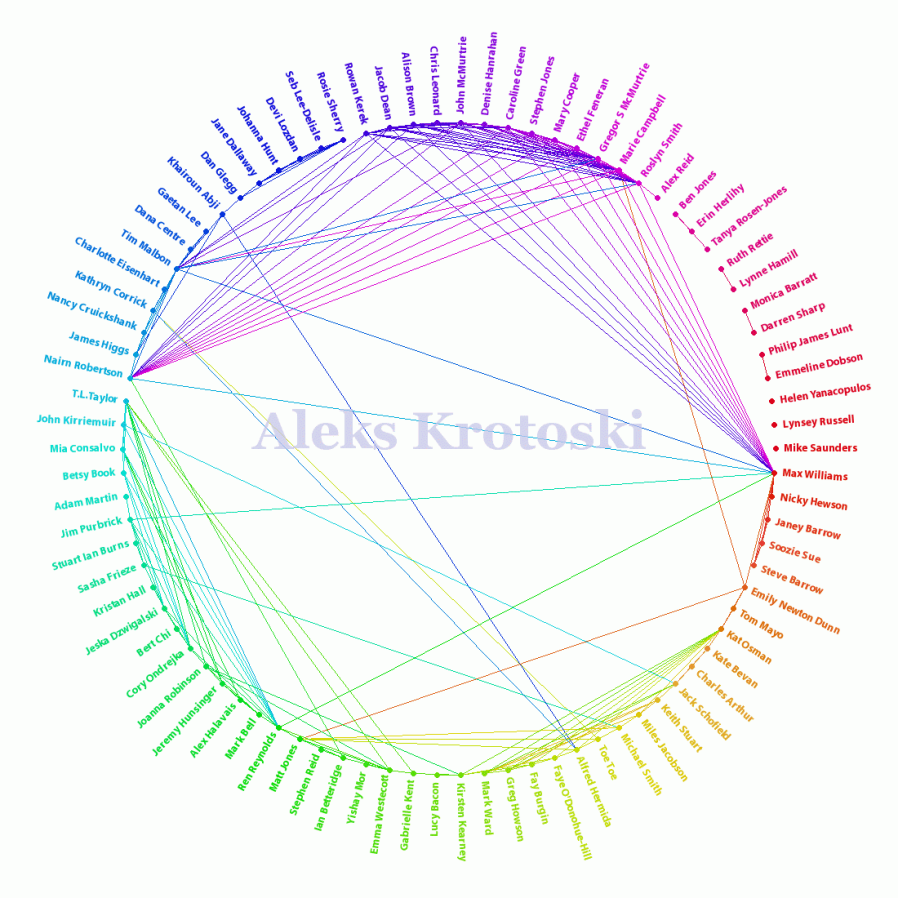

Figure 7: Facebook Friend Wheel visualization of relationships between a user's friends [12]

2.1.1.2 *Glyphs.* Glyphs are graphical objects, or icons, used to visually represent multivariate data [2]. Differences in distinguishing features of the glyphs highlight differences in the data points they represent. Glyphs exploit many of the preattentive features discussed in Section 2.1 such as size, color, texture, and orientation to convey information. There is no restriction on the graphical objects that can be used as glyphs. They could be simple shapes like circles or squares, or they could be more complex images depending on the data to be represented. Figure 8 shows a simple application of glyphs overlaid on part of a network map.

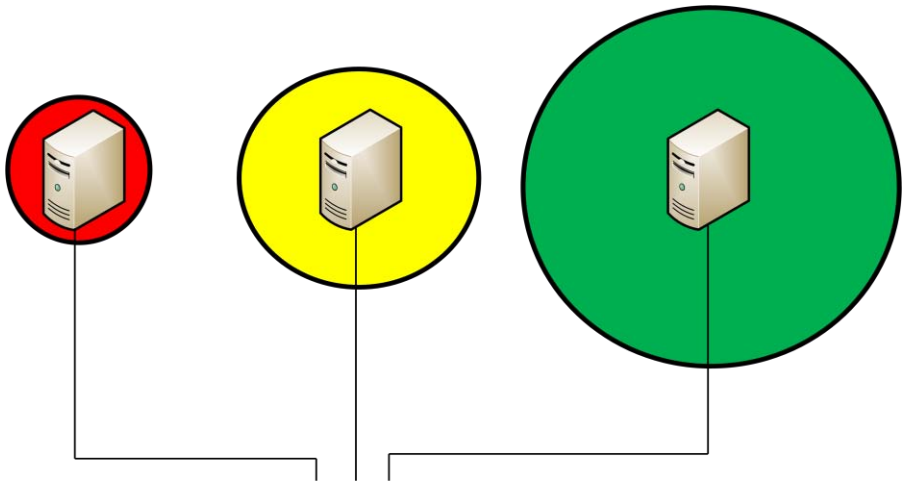

Figure 8: Example of simple glyphs applied to a network map (size represents amount of data processed, red = encrypted data, green = unencrypted data, and yellow = both types of data)

In the example in Figure 8, simple circles are used to quickly convey additional information about servers on a network. Color is used to indicate the type of information being processed by each node: red for encrypted data only, yellow for both encrypted and unencrypted data, and green for unencrypted data only. The size of the circles represents the amount of traffic each server is processing. While it is impossible to determine precisely how much data each server is processing without explicitly including that information, the glyphs make it easy to see at a glance which server has the highest workload. An administrator could be alerted to a potential issue simply through a change in the typical appearance of the glyphs. A sudden increase in the size of a glyph monitoring the number of connections to a server might indicate malicious activity without the administrator requiring the precise number of connections before taking certain actions.

2.1.1.3 *Treemaps.* Treemaps were developed in the early 1990s as a space-filling approach to representing multidimensional, hierarchical data [5] [13] [14]. They were developed as a way of visually representing directory tree structures within a confined space rather than producing extremely large, unmanageable node-link graphs. The basic concept is fairly simple: divide the display space into rectangles alternating between horizontal and vertical divisions with each new directory level. Just as each new division can represent an additional level in a directory tree, the divisions can also be used to represent a different dimension in a dataset. As a result, treemaps have become increasingly popular and have been applied to numerous datasets [14]. Treemaps can also encode additional information through other attributes of the rectangles within the display such as size, color and texture. The treemap in Figure 9 illustrates the concept using simple firewall log information. In Figure 9, the display space is divided horizontally to separate two types of systems behind a firewall: financial systems and mail servers. Each section is further subdivided to represent each individual system within those categories. The size of the boxes indicated the amount of traffic passed or blocked by the firewall, and color immediately distinguishes between traffic that was allowed or denied.

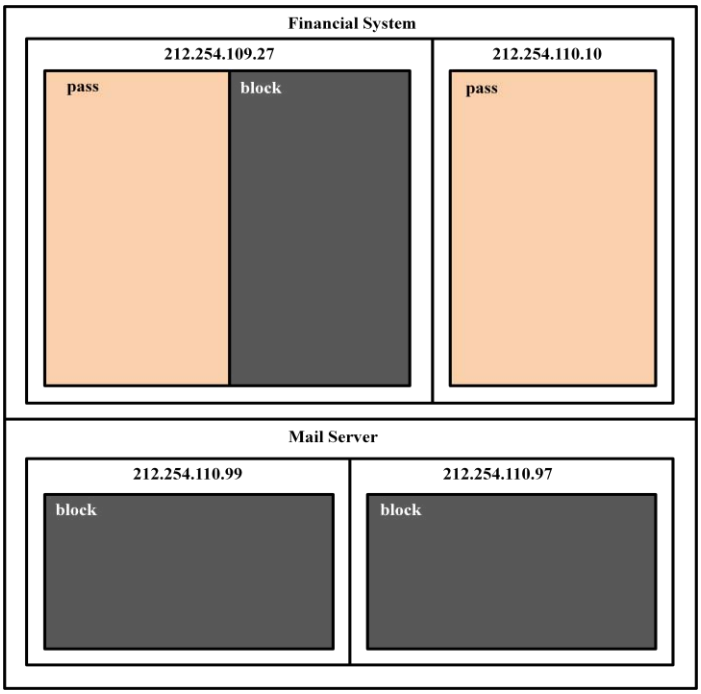

Figure 9: A treemap representing simple firewall log information [5]

2.1.2. Network Visualization

Network management can be a complicated endeavor and is an area that benefits greatly from the use of visualization. Network administrators are often responsible for managing networks with hundreds of workstations, servers, routers, switches, and peripherals. Different physical media, connection types, software applications, and other related network components further complicate matters. Add a healthy dose of security considerations, and network administrators have to keep track of a great deal of information. They need tools that can present as much of this information as concisely as possible and in a way that is easily understood.

Even a simple example network can effectively illustrate the benefits of visualization in networking. Basic network topology can be represented as graph data.

Consider the information provided in Table 1, which provides the connection information for an 8-node network. The numbers preceded by dots labeling the matrix represent the last octet of the nodes' IP addresses. An entry of "1" represents a connection between two nodes.

Table 1: Adjacency Matrix for an 8-node network [4]

	.1	.2	.3	.4	.5	.6	.7	.8
.1	0	1	0	0	0	0	0	0
.2	1	0	1	0	0	0	0	0
.3	0	1	0	1	1	1	1	1
.4	0	0	1	0	0	0	0	0
.5	0	0	1	0	0	0	0	0
.6	0	0	1	0	0	0	0	0
.7	0	0	1	0	0	0	0	0
.8	0	0	1	0	0	0	0	0

Though it can quickly be determined if there is a connection between any two specific nodes, it takes quite a bit of processing to track all of the connections from the matrix and develop a clear understanding of the network topology. Also, considering the effort required to draw a clear mental picture of the network from the table for an 8-node network, imagine how tedious and difficult it would be if there were hundreds of nodes.

Since we are, in fact, trying to develop a clear picture of the layout of the network, let us use actual pictures. The information presented in Table 2.1 can be represented visually using a simple link graph. Using circles containing the labels to represent nodes and line segments to represent connections, the same network is depicted in Figure 10.

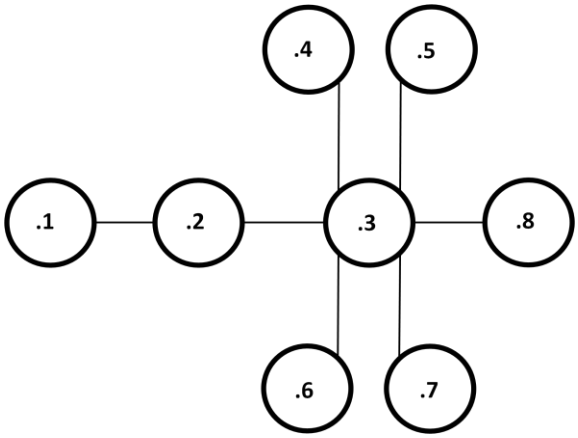

Figure 10: Graphical visualization of 8-node network in Table 2.1 [4]

With this simple visualization, the viewer has an instant overview of the network. The same basic information is conveyed, but the visualization allows the viewer to process more of the data in parallel. A complete picture of how the nodes are connected is presented, and individuals with even limited networking experience could infer information about some of the nodes based on the layout.

Of course, the visualization in Figure 2.10 does not provide all of the network information an administrator would require. There is no way to differentiate between nodes in Figure 10. Without some additional information, this visual depiction of a network is not much practical use. Examine Figure 11, which represents the same network using more distinguishable icons for the network nodes.

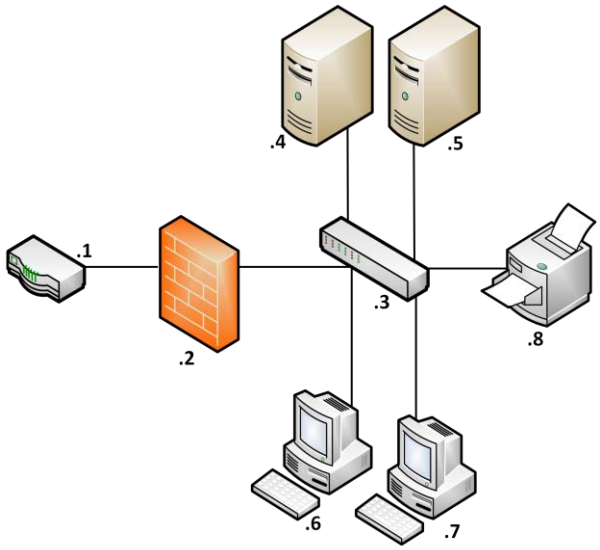

Figure 11: Another graphical visualization of 8-node network in Table 2.1

Here we can begin to see the true advantages of network visualization. Using essentially the same amount of space, Figure 11 conveys much more information than Figure 10. Now, most anyone with basic networking knowledge could identify the various networking components, including differentiating between servers, workstations, and peripherals. While the visualization could be enhanced further, the viewer is instantly presented with a clear picture of both the network topology and the basic role of each of its components. Even using such a small example, the benefits of visualization in network management are easily recognizable.

Now consider the additional information that could be encoded into Figure 11 without requiring much, or any, additional space or cluttering the visualization. Line thickness could be changed to indicate bandwidth or physical media differences. Icons could be colored to indicate status on the network. The workstation icons could be

relabeled to represents entire subnets. If any node(s) on a given subnet went down, the icon could change to yellow prompting an administrator to click the icon to display a detailed map of that subnet and identify the issue. The icon could turn red if connectivity to the entire subnet was lost.

Visualizations, for better or worse, are truly only limited by imagination. One thing is for certain though: visualization plays a key role in delivering large amounts of data to humans quickly and in ways that are more naturally processed and understood. This is extremely important in many critical areas today such as data processing, network management, and situation awareness.

2.1.3. Situation Awareness

In the most general terms, situation awareness is knowing what is going on around you. Much work has been done to more precisely define situation awareness and the factors that affect it either in general or in a context specific to a given domain. One of the most widely accepted definitions for situation awareness comes from Dr. Mica Endsley [15] as, "the perception of elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future." The Endsley Model [16], shown in Figure 12, illustrates Endsley's proposed components of situation awareness: perception, comprehension, and projection; the relationship of situation awareness to decision making (most importantly that they are separate from one another); and the factors that influence and are influenced by situation awareness.

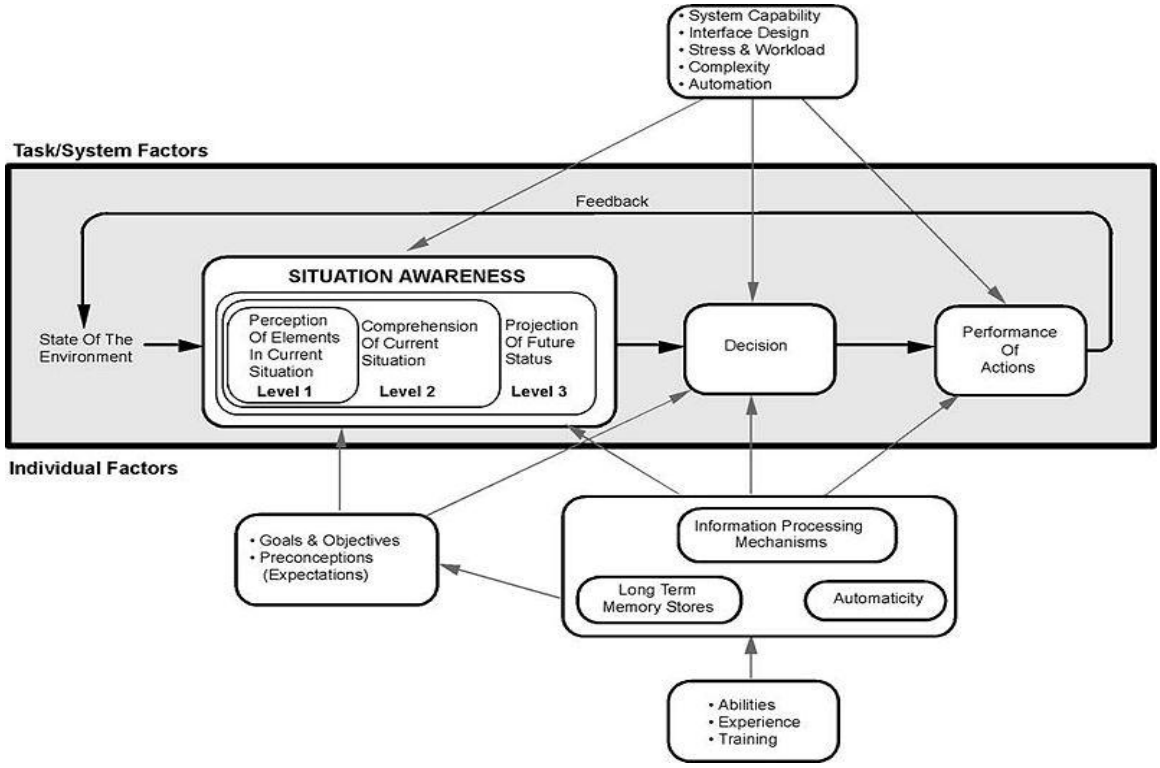

Figure 12: The Endsley Model of Situation Awareness [16]

Visual network management and intrusion detection tools follow this process. These tools collect raw data from the network and present it to human operators to enable understanding of the network's current state. Experience allows operators to make inferences regarding causes and potential results of given events. The situation awareness gained through the tools allows operators to make decisions and take actions that affect the environment and then feed back into the situation awareness cycle.

Much work has been done to adapt the concepts and definitions of situation awareness research to the cyber domain [17] [18]. Cyberspace means different things to different people, and as a result, so does cyber situation awareness. To some it may mean knowledge about the status of network devices while others may be concerned only with

malicious activity. An argument could certainly be made that true cyber situation awareness would consist of both and more. Regardless of how it is defined in any one instance, the amount of information being transmitted through cyberspace will likely continue to rapidly grow, and visualization seems to be the best approach for allowing us to effectively sift through it all.

Network visualization tools at their core aim to increase our situation awareness. Whether it is designed to ease the burden of management or assist in decision making, visualization tools are used to enhance our ability to more completely understand the environment they are deployed in. As the focus on cyber operations intensifies and the amount of data that is processed increases, our reliance on visualization to help us make sense of it all increases too. One example that clearly illustrates this is Cybercraft [19].

2.1.3.1 *Cybercraft.* The Cybercraft project was started in 2005 by the Air Force Research Laboratory as a potential approach to defending Air Force networks [19]. It is meant to fill the gaps in network intrusion detection and network defense that may be identified at points in the future and that are not addressed by the system or systems in use at that time. The natural analogy is the creation of a vehicle (a Cybercraft) that operates in the cyber domain to carry out specific missions similarly to how air and spacecraft operate in the air and space domains. When an objective is identified, a Cybercraft payload can be developed and deployed within the Cybercraft framework to address the issue. Cybercraft could potentially be deployed to process and deliver a wide spectrum of information. One payload may be designed to relay information from low-level sensors to detect potential insider threats [18] while another may be used to evaluate

and update IPsec policies throughout the network. As Cybercraft carry out their missions, it is necessary for them to provide potentially massive volumes of data back to human operators in a way that is readily accessible and understandable.

A visualization tool is specified as one of the required Cybercraft support products to meet this requirement and help achieve enhanced situation awareness [19]. This visualization requirement supports the overarching concept that visualization is an accepted, flexible, and powerful approach to complex problems that require processing large datasets, especially when they require some degree of human interpretation. It follows then that visualization, specifically through the benefits it provides in network management, is perhaps the ideal approach to address the complexity of IPsec management.

2.2. IPsec

IPsec is a powerful but complicated protocol suite designed to provide cryptographically-based security for networks using the Internet Protocol. The original Internet Protocol (IPv4) was designed with a 32-bit address space under the assumption that this would provide more addresses than would likely ever be used. Likewise, IP was not designed with security built in because the security issues we face today were not a consideration when computer networking was being developed. When IPv6 was developed to address the ever-shrinking IPv4 available address space due to the unexpected popularity of the Internet, the decision was made to also address the lack of security. However, it was clear that there would not be an immediate transition to IPv6

and that security was an immediate issue. As a result, IPsec, which is built in to IPv6, was designed to be compatible with IPv4 [20].

IPsec relies on several protocols and cryptographic standards to provide security to individual IP packets based on rules and policies set by the user. The following sections provide the basic information necessary to understand IPsec and how it provides network-layer security. More detailed information on sections 2.2.2 through 2.2.6 is provided in various IPsec-related RFCs [1] [21] [22] [23] and references [20] [24] [25] [26].

2.2.1. Protocols

The Internet Protocol is part of the TCP/IP protocol suite. This can be visualized as a stack of layers as shown in Figure 13, where each layer is responsible for a specific piece of the communication process.

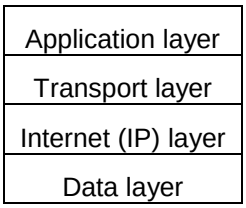

Figure 13: The TCP/IP protocol stack [24]

When outbound communication occurs, each layer from top to bottom appends a header to the data to be transmitted and passes the information to the layer below. When inbound communication is received, the process is reversed, and data is passed up the

stack with each layer stripping out its header. IPsec operates at the Internet (network) layer and makes use of the IP header, illustrated in Figure 14.

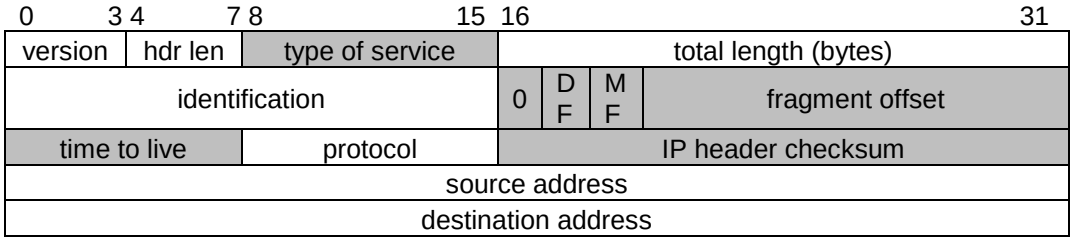

Figure 14: The IP header format [26]

The shaded fields are mutable (fields that may be changed by intermediate nodes as the packet travels from source to destination). Due to the dynamic nature of these fields, they must be treated specially when calculating values for authentication as discussed below.

IPsec relies on two protocols, the Authentication Header (AH) and Encapsulating Security Payload (ESP), to provide its security services of integrity, authentication, confidentiality, and anti-replay protection. These protocols can be used separately or in combination to provide varying level of security under different circumstances.

2.2.1.1 *Authentication Header* [21]. The Authentication Header (AH) provides connectionless integrity, data origin authentication, and replay protection. Connectionless integrity means that no tampering has occurred to the message in transit. It is connectionless because no attempt is made at the IP layer to ensure proper delivery of information. That is left to the transport layer or the originating application. Data origin authentication guarantees the message has not been spoofed and was actually sent

by who it appears to have been sent by. Replay protection ensures messages are not delivered multiple times or out of order through the use of sequence numbers. The sender must implement this capability, but it is optional for the receiver to make use of it.

The Authentication Header protects an IP packet by calculating an integrity check value (ICV) over the payload and the non-mutable fields of the IP header. The mutable fields must be zeroed out before the ICV can be calculated. The ICV is stored in the authentication data field of the AH header, shown in Figure 15, and inserted into the IP packet.

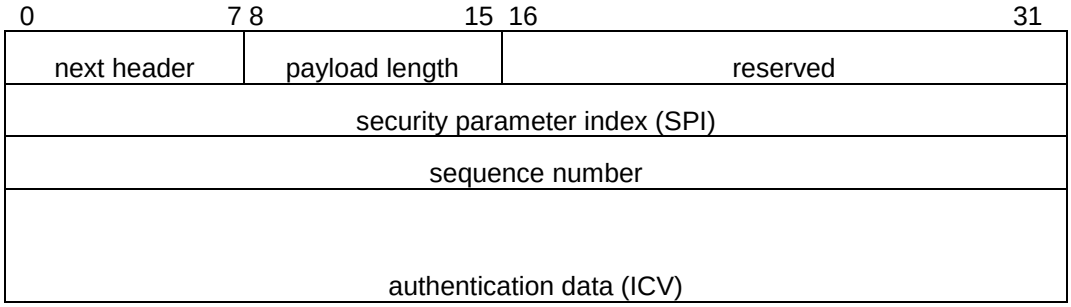

Figure 15: The AH header format [26]

The IPsec mode being used determines where the AH header is placed within the IP packet. IPsec modes are discussed below. The ICV can be calculated by any MAC provided in an implementation, but all implementations are required to support at a minimum HMAC-MD5-96 and HMAC-SHA1-96.

2.2.1.2 *Encapsulating Security Payload* [22]. The Encapsulating Security Payload (ESP) provides confidentiality, limited traffic flow analysis protection, connectionless integrity, data origin authentication, and replay protection.

Confidentiality means that any captured, intercepted, or otherwise viewed data cannot be understood by anyone other than the intended recipient. Limited traffic flow analysis means that to some extent an eavesdropper cannot discern who is communicating with whom or gather precise information regarding the nature of the communications.

The Encapsulating Security Payload (ESP) provides its authentication, integrity, and replay protection services in much the same way AH does. The data used for authentication and the placement of the authentication data in the ESP packet differs slightly from AH though. Figure 16 shows the format of the ESP packet.

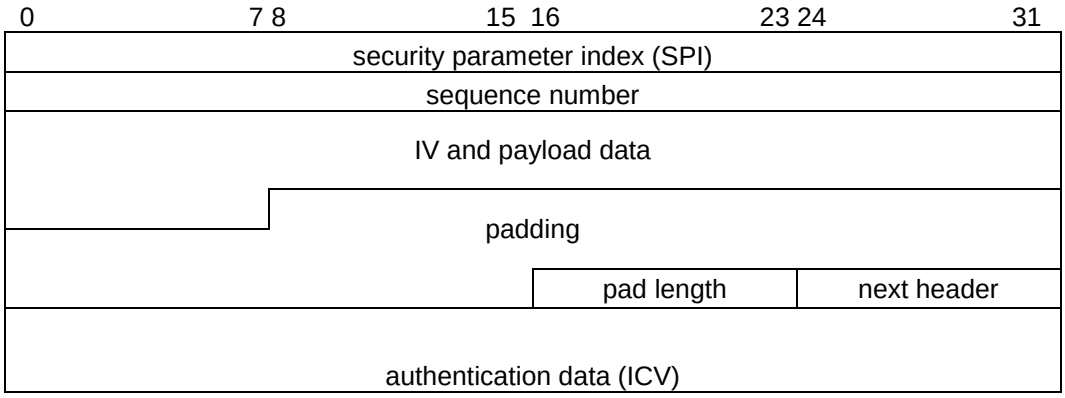

Figure 16: The ESP packet format [26]

The ICV is calculated over the ESP packet, excluding the authentication data field where the ICV is stored. When considering how the ESP packet is integrated into the IP packet, the following terminology is typically used:

- ESP header: the SPI and sequence number fields
- ESP trailer: the padding, pad length, and next header fields
- ESP authentication data: the ICV

2.2.2. Modes

IPsec has two modes of operation; transport mode and tunnel mode. The security protocols discussed above, AH and ESP, can both operate in either mode. The two modes differ in how they encapsulate data, which dictates the situations they are typically used in.

2.2.2.1 Transport Mode. Transport Mode is typically used for communication between two fixed hosts where each host is the final destination of the protected communication. Specifically, it cannot be used to connect a host or range of hosts to a network. Figure 17 shows how data is encapsulated in Transport Mode. The single IP header explains why each host must be the final destination for communication as opposed to tunnel mode described below.

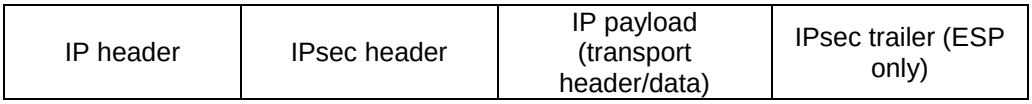

Figure 17: Transport Mode encapsulation [26]

As stated earlier, both the Authentication Header and Encapsulating Security Payload protocols can operate in Transport Mode. Figure 18 illustrates where the AH header is added to the IP packet in Transport Mode. Here, the AH header is inserted between the IP header and the payload of the IP packet, which includes the transport layer header and the message data. Having only half of the IP header identified as being authenticated illustrates the fact that only the non-mutable fields in the header are used for authentication. The mutable fields are zeroed out before the ICV is calculated.

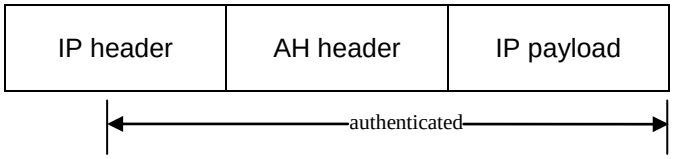

Figure 18: AH Transport Mode encapsulation [26]

With ESP, the ESP packet is broken down into separate parts to provide its protection services. Figure 19 shows ESP encapsulation in Transport Mode. The ESP header cannot be encrypted because the SPI field is one of the fields used to determine which IPsec policy applies to the packet. If the ESP header was encrypted, the receiver would not be able to process the packet.

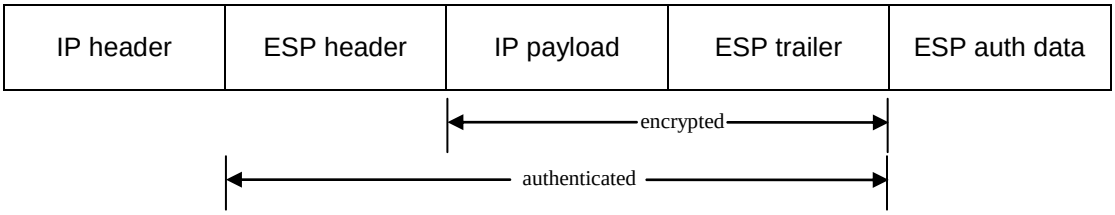

Figure 19: ESP Transport Mode encapsulation [26]

2.2.2.2 *Tunnel Mode*. Tunnel Mode is used to secure communications between two networks or a remote host and a network. In contrast to Transport Mode, Tunnel Mode is used when the endpoints of the protected communication are not necessarily the final destination. For example, a Virtual Private Network (VPN) may exist between two networks where any communication from a host on one network

destined for any host on the other is encrypted between a router on each network. The routers themselves are not the final destinations of the communication, but the data is only protected on the path between the two routers.

Figure 20 shows how Tunnel Mode allows for secure communication between to endpoints (of a tunnel) that are not the final destination of the communication. The outer IP header is the tunnel endpoint. Once the packet reaches this endpoint, the IPsec protection ends, the outer IP header and IPsec header are stripped away, and the final destination contained in the inner IP header is revealed.

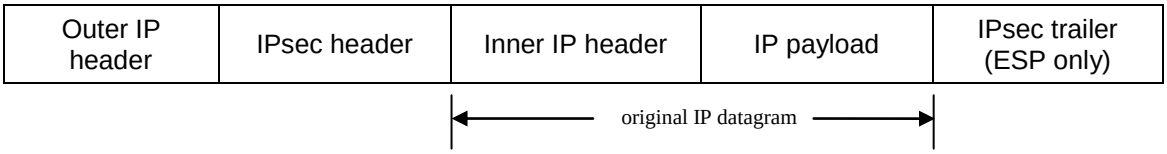

Figure 20: Tunnel Mode encapsulation [26]

AH authentication in Tunnel Mode is similar to that of Transport Mode except that the AH header, along with the outer IP header, are prepended to the original (inner) IP header, and the fields of the outer IP header are used in calculating the ICV. Figure 21 illustrates AH encapsulation in Tunnel Mode.

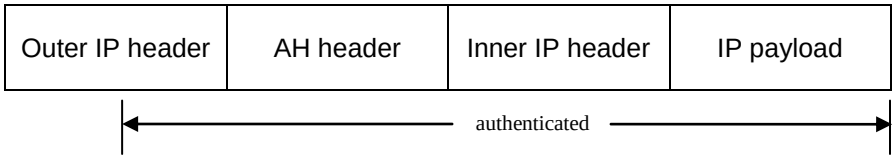

Figure 21: AH Tunnel Mode encapsulation [26]

The same modification occurs with ESP in Tunnel Mode where the original IP datagram becomes the payload, making the original IP header the inner IP header, and the tunnel endpoint is added as the outer IP header. Figure 22 shows ESP encapsulation in Tunnel Mode.

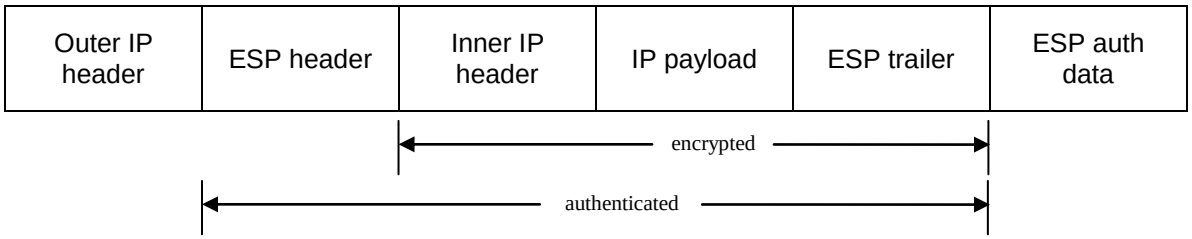

Figure 22: ESP Tunnel Mode encapsulation [26]

To further complicate things, IPsec also supports nested tunnels with regards to both protocols and modes as well. For example, in the scenario above where two networks are connected by a VPN that encrypts data between two routers, the hosts could have an IPsec policy that provides authentication between them as well. This could be accomplished by a pair of AH Transport Mode security associations between the two hosts and a pair of ESP Tunnel Mode security associations between the routers as depicted in Figure 23.

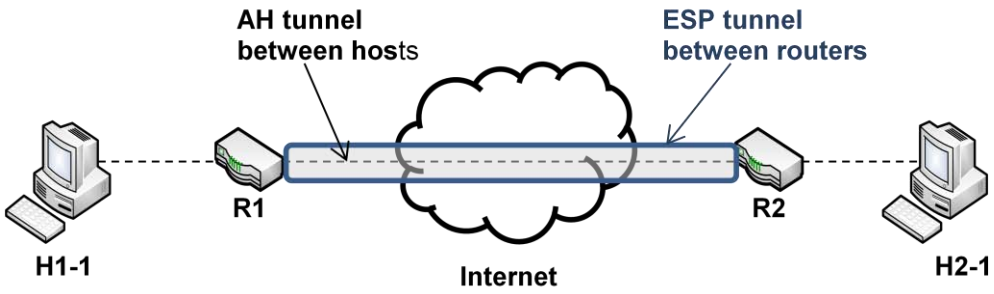

Figure 23: Example of a nested tunnel [25]

2.2.3. Cryptography

Cryptography is what enables IPsec to provide security services to packets of data. Both authentication and encryption services rely on cryptographic algorithms [27], albeit in different ways, to provide data protection.

2.2.3.1 *Authentication Algorithms*. Authentication algorithms are based on one-way hashes. A one-way hash takes a message as its input and computes a value in a way that meets several criteria:

- given the hash, it is computationally infeasible to recover the original message
- it is computationally infeasible to find two different messages that hash to the same value
- given a message and its hash, it is computationally infeasible to find another message with the same hash value

Typically, the hash would be transmitted with the original message, and the receiver could then compute the hash and compare it to the one received to see if the message had been changed. However, someone intercepting traffic could simply change

the message and recompute the hash. A way is needed to ensure not only that the message has not been tampered with, but that it actually comes from who it appears to have been sent by. This requires a keyed hash, called a message authentication code (MAC), that introduces information from a secret key into the hash value. Using this method, a secret key known only to the sender and receiver ensures that if the MAC sent can be recomputed by the receiver, the message has not been changed and is from the originator. IPsec mandates that all implementations must support HMAC-SHA-1 [27], which is based on the SHA-1 [28] algorithm, for AH authentication services. HMAC-SHA1-96 [27] is the mandatory-to-support algorithm for ESP authentication. Other algorithms can be supported by either protocol.

2.2.3.2 *Encryption Algorithms*. Unlike authentication algorithms, which protect the integrity of a message and its origin, encryption algorithms protect the data itself. Using a secret key, encryption algorithms transform original messages, or plaintext, to an encrypted form called ciphertext. The encryption algorithms used by ESP are block ciphers that operate in Cipher Block Chaining Mode (CBC). This means that each block of text to be encrypted is XOR'd with the encrypted text resulting from the previous block's encryption. The first block of text has no previous text for the operation, so it is XOR's with an initialization vector (IV). IPsec generates a random IV for each packet. Currently, the only encryption algorithm that must be supported by an IPsec implementation is TripleDES-CBC [29], but it will likely be replaced by AES-CBC with 128-bit keys [30].

Since authentication and encryption are both optional services with ESP, all IPsec implementations must also support a "NULL" for both authentication and encryption. Setting either of these to "NULL" indicates that the service specified is not being provided. However, if ESP is being used, one service must be selected, so authentication and encryption cannot be set to "NULL" at the same time.

2.2.4. Security Associations

Security associations (SAs) are the mechanisms that allow IPsec-protected communication between two entities. SAs encapsulate the shared state of these two entities, or endpoints, which consists of the information that enables IPsec protection. This information includes which IPSec protocols and modes are used, which authentication and encryption algorithms are used and their keys, the lifetime of the SA, and an identifying number called the security parameters index (SPI). IPsec relies on the SPI to ensure the right SA is chosen when processing packets. SAs rely on the security policy database (SPD) to be successfully negotiated. The SPI and SPD are discussed below, followed by a more thorough exploration of security associations.

2.2.4.1 *Security Parameters Index.* As identified in Section 2.2.1, both the Authentication Header and Encapsulating Security Payload headers have an SPI field. The SPI is a 32-bit value assigned by the destination host when an SA is first negotiated. The SPI is used in conjunction with the destination address and protocol (AH or ESP) fields to uniquely identify the appropriate SA to use.

2.2.4.2 *Security Policy Database*. Security policy is the heart of IPsec. Security policies identify, by source/destination address, source/destination port, and

transport protocol, what traffic requires IPsec protection. These policies, or rules, specify what actions should be taken for traffic that matches these rules, and when IPsec protection should be afforded to a packet, the SPD specifies how it is accomplished. SPD rules can result in one of three possible actions:

- Drop the packet: an SPD rule can specify that traffic meeting certain criteria be dropped. Additionally, packets are dropped if there is no SPD rule to match the traffic.
- Process without protection: rules can dictate that specific traffic is processed in the clear.
- Apply IPsec protection: when IPsec protection is required the SPD specifies the protocol, authentication and encryption algorithms, and mode to be used.

The SPD serves different roles for processing outbound traffic than it does for inbound traffic. For our purposes, it is simple enough to consider the SPD as a table listing the rules for a given host and there being one for outbound traffic and one for inbound traffic. Rules in the SPD are chosen based on selectors, which can be source and destination address, source and destination port, or the transport protocol used.

Consider the scenario depicted in Figure 24, where two networks, N1 and N2, are joined by two security gateways, SG1 and SG2. Table 2 shows a possible SPD for one of the security gateways in Figure 24.

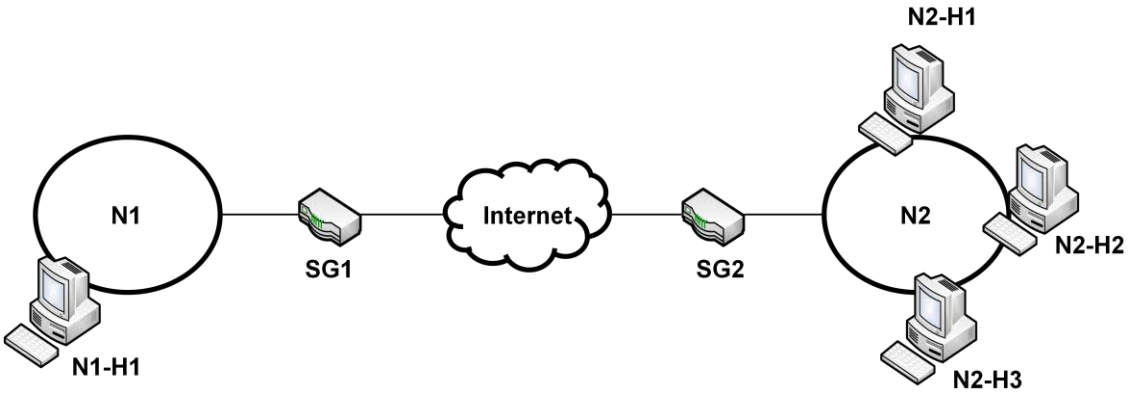

Figure 24: Two networks connected through security gateways [24]

Table 2: Sample Security Policy Database for a security gateway [24]

Rule	Src Addr	Dest Addr	Src Port	Dest Port	Prot	Action	IPsec Hdr	Enc Alg	Auth Alg	Mode
1	SG1	SG2	500	500	Any	Accept	/	/	/	/
2	SG1	SG2	Any	Any	Any	IPsec	AH	/	HMAC-SHA-1	Tunnel
3	N1-H1	N2	Any	Any	Any	IPsec	ESP	AES	HMAC-SHA-1	Tunnel
4	N1	N2	Any	Any	Any	IPsec	ESP	3DES	HMAC-SHA-1	Tunnel

Since SPD rules are required for both outgoing and incoming traffic, the source and destination entries show that this could be either the outbound SPD for SG1 or the inbound SPD for SG2. The rules in Table 2.2 illustrate some of the ways to identify traffic requiring protection and how that protection may be applied.

- Rule 1: allows any traffic from SG1 on port 500 destined for SG2 on port 500 to be sent with no protection
- Rule 2: requires AH authentication on all other traffic from SG1 to SG2, regardless of port, using HMAC-SHA-1

- Rule 3: requires any traffic originating from H1-1 destined for any host on N2 be authenticated using ESP HMAC-SHA-1 and encrypted using AES
- Rule 4: specifies that any other traffic originating from N1 and destined for N2 be authenticated using ESP HMAC-SHA-1 and encrypted using 3DES

There are two aspects of the SPD that are important to note. The first is that the order of the rules is important. Rule 1 is more specific than Rule 2. If Rules 1 and 2 in Table 2.2 were switched, Rule 1 would never be used because the traffic specified in Rule 1 matches the more general criteria specified in Rule 2. Second, as mentioned above, if traffic is encountered that does not match any rule in the SPD, it should be dropped. This may not be desirable and would require a default rule to catch traffic not already specified in the SPD. Based on the fact that order matters, it is imperative that if such a rule is used, it should be placed at the end of the SPD.

2.2.4.3 *Security Associations*. SAs are simplex communication channels. This means that any host that is the endpoint of an IPsec-protected channel will typically have a pair of SAs: one for inbound traffic and one for outbound traffic. Furthermore, an SA can only accommodate one IPsec protocol. If a channel uses both AH and ESP for security services, each endpoint will have four SAs.

As mentioned above, the SPI, destination address, and IPsec protocol are used to uniquely identify an SA. Multiple SAs can be required at once as in the case where a channel is protected by both IPsec protocols. When this happens, the SAs are combined into groups called SA bundles. Additionally, a rule in the SPD can spawn either a single SA or multiple SAs when one of its selectors identifies more than a single entity.

Consider rule 3 in Table 2.2 above. It addresses traffic from H1 on N1 destined for any host on N2. This could result in a single SA covering traffic from H1 to N2 as shown in Table 2.3 or multiple SAs from H1 to the individual hosts on N2 as shown in Table 2.4.

Table 3: Single SA generated from SPD rule 3 in Table 2.2 [24]

SA	Src Addr	Dest Addr	Src Port	Dest Port	Prot	IPsec Hdr	Enc Alg	Auth Alg	Mode
1	H1-1	Any	Any	Any	Any	ESP	AES	HMAC-SHA-1	Tunnel

Table 4: Multiple SAs generated from SPD rule 3 in Table 2.2 [24]

SA	Src Addr	Dest Addr	Src Port	Dest Port	Prot	IPsec Hdr	Enc Alg	Auth Alg	Mode
1	H1-1	H2-1	Any	Any	Any	ESP	AES	HMAC-SHA-1	Tunnel
2	H1-1	H2-2	Any	Any	Any	ESP	AES	HMAC-SHA-1	Tunnel
3	H1-1	H2-3	Any	Any	Any	ESP	AES	HMAC-SHA-1	Tunnel

SAs are stored in a construct called the Security Association Database (SADB). Similar to the SPD, the SADB is a notional construct, and its details are implementation specific. The SADB is used in both outbound and inbound processing. When outbound traffic requires protection, IPsec searches the SADB for the appropriate SA to determine what parameters need to be applied. If an appropriate SA cannot be found, IPsec will attempt to negotiate one. For inbound traffic, IPsec searches the SADB using the SPI, destination address, and protocol for the SA containing the necessary parameters to authenticate and/or decrypt the packet.

Finally, there are two ways an SA can be created: manual keying and the Internet Key Exchange protocol. Manual keying is fairly simple and involves administrators on

both sides of the communication channel to agree on the parameters of the SA offline. While adequate for testing purposes, manual keying does not scale well, and SAs created using manual keying do not expire. Long-lived cryptographic keys present a known security risk. The other option is to use a key management protocol such as the Internet Key Exchange protocol (IKE) [23]. IKE can automatically negotiate SAs based on existing policies, can negotiate new SAs when existing ones are about to expire, and can provide perfect forward secrecy. IKE is discussed in more detail in Section 2.2.5.

2.2.5. Key Exchange

IPsec uses the Internet Key Exchange (IKE) protocol for negotiating IPsec Security Associations (discussed below) and key management. Through IKE, peers can authenticate each other and establish a secure channel for exchanging keying material, obtained through a Diffie-Hellman exchange [26], for authentication and encryption services provided by IPsec security protocols.

A Diffie-Hellman exchange allows two peers to exchange public values over an unprotected network and combine them with private values known only to each individual peer to arrive at a shared secret value. This shared secret can then be used as the key for an encryption algorithm to protect communications between the two peers. Figure 25 illustrates the Diffie-Hellman exchange for key generation.

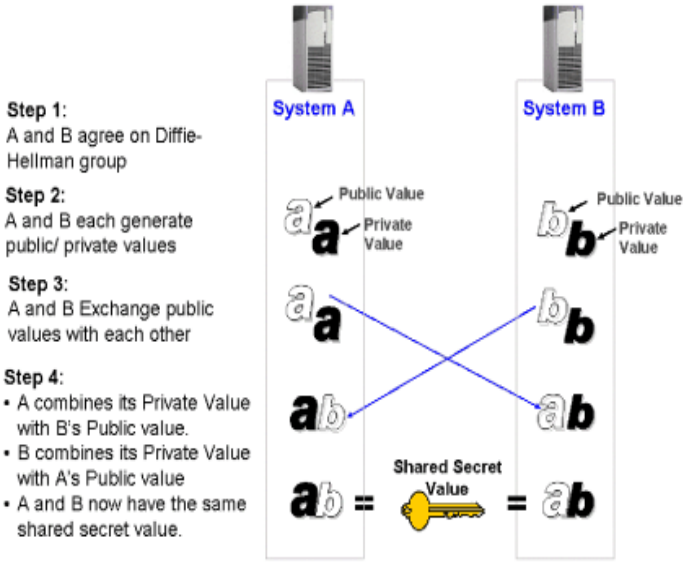

Figure 25: Diffie-Hellman key exchange [31]

IKE is a hybrid protocol, pulling different parts from three other key-exchange protocols [26]: the Internet Security Association and Key Management Protocol (ISAKMP) [37], the Oakley Key Determination protocol (OAKLEY) [38], and the SKEME protocol [39]. The elements that IKE borrows from each of these protocols are described below.

- ISAKMP: forms the base of IKE. ISAKMP is a framework that provides the foundation for protocols that need to do key exchange and set up security associations. IKE uses several concepts, exchanges, payloads, and message types from ISAKMP.
- OAKLEY: specifies several methods for using the Diffie-Hellman algorithm for secure key exchange over an unsecure channel. IKE borrows the concept of using different modes to achieve a secure key exchange as well as several fixed groups used for the Diffie-Hellman exchange.

- SKEME: defines an authenticated key exchange using public-key cryptography, which can be combined with a Diffie-Hellman exchange for perfect forward secrecy. IKE uses this as one of its authentication methods and also uses a technique from SKEME for fast key renewal using nonces.

IKE uses a two-phase exchange to negotiate an IPsec security association (SA). The first phase has two modes, main mode and aggressive mode, and is for establishing a secure, authenticated channel between two peers as well as authenticated keying material for protecting communication over this channel. Main mode is considered more secure as it protects the identity of the sender and receiver while aggressive mode requires fewer messages and therefore takes up less bandwidth. There are four ways of authenticating peers during phase one: shared secret, digital signature, public key encryption, and revised public key encryption. Each method has its advantages and disadvantages stemming from complexity and the overhead caused by the required calculations.

The second phase, using a quick mode exchange, is for negotiating the IPsec SA. Quick mode is so named because, if perfect forward secrecy is not required, SAs can be negotiated quickly without requiring a Diffie-Hellman calculation. Additionally, it is possible to negotiate several SAs in a single quick mode exchange by including multiple SA payloads in the exchange.

2.3. Summary

This chapter introduced the background information behind data visualization, specifically the impact of visualization on network management and situation awareness. This chapter also explored the many components and operations of IPsec and how IPsec

could be used to protect local area network traffic. The fusion of network visualization and IPsec technologies and concepts is the foundation of this research effort and forms the basis of the work shown in the chapters that follow. To simplify IPsec management, we will explore ways to streamline the IPsec implementation presented to an administrator, visualize IPsec rules, and develop functionality through visualization that does not exist in today's IPsec management tools.

III. Approaches to Simplifying IPsec Management

The information presented in Chapter II provides the building blocks required to apply visualization techniques to IPsec management. There are many areas that would require attention to fully realize a tool for IPsec management. This chapter identifies the specific areas this thesis addresses and outlines the methodology to be used in attacking each area.

Addressing all of the factors involved in developing a fully-functional IPsec management tool constitutes a complete software development effort and is outside the scope of this document. Issues such as security and performance are important considerations, but they have little impact on the applicability of an approach towards visualization or simplifying management. Scalability is a concern in developing both network management tools and data visualizations due to the amount of information that needs to be presented to the user. However, an initial approach must be developed and validated before scalability issues can feasibly be addressed.

This thesis focuses on two areas in attempting to simplify IPsec management: simplifying the IPsec implementation presented to the administrator and visualizing IPsec rules deployed throughout a network. The former has a direct impact on the complexity of both IPsec management and the visualization of IPsec rules, and the latter is the crux of this research effort, attempting to provide administrators with an efficient method for managing a potentially complex dataset. Additionally, functionality for managing an IPsec deployment and issues with visualizing the necessary data are explored. Together these areas form the foundation for developing a comprehensive IPsec management tool.

3.1. Simplifying IPsec

The first step in simplifying IPsec management is simplifying IPsec itself. Although modifying the actual IPsec protocol suite is outside the scope of this document, streamlining how it might be implemented and presented to an administrator is not. Simplifying the IPsec implementation not only reduces management complexity for the administrator, it reduces the complexity inherent in any visualization used to represent it. The fewer components the system has, the fewer distinct pieces of information need to be represented in a visualization. However, this thesis does not suggest that the streamlining decisions presented here are the optimum configuration for an IPsec implementation. This thesis simply recognizes that such suggestions exist in the literature and evaluates the effects making such decisions have on administration and visualization.

3.1.1. Problem Definition

The problem is that there are an excessive number of options and resulting possible configurations available to administrators when implementing IPsec rules between peers. As a result, an administrator must make many decisions for each IPsec rule implemented and any visualization chosen to represent IPsec rules must reflect a large number of distinct data points.

3.1.1.1 *Goals and Hypotheses*. The goal is to evaluate the effects of simplifying IPsec administration by reducing the number of options available to administrators. This will reduce the number of possible configurations and decisions to be made when implementing IPsec rules and simplify any visualization chosen to represent those rules. The hypothesis is that some parts and operations of IPsec can be

eliminated or set to default values that result in simplified administration without negatively impacting the effectiveness of IPsec.

3.1.2. Approach

IPsec is complex. There is some degree of overlap in some features regarding the types of services provided and how data is protected. When nothing but authentication between two machines is desired, an administrator has four options to choose from: AH in transport mode, AH in tunnel mode, ESP in transport mode, and ESP in tunnel mode. When both authentication and encryption are required, the number of possible configurations increases. ESP could be used to provide both services, or a combination of ESP and AH could be used, again in either transport or tunnel mode, and now with either a single IPsec rule or multiple rules creating nested tunnels.

Each of the major components of IPsec is evaluated to identify areas where the implementation can be streamlined to reduce complexity. Whenever possible without significantly degrading the security afforded by IPsec, options are eliminated or set to default values to reduce the number of decisions and administrator must make. The goal here is to show that streamlining the IPsec implementation reduces the complexity of both management and visualization of IPsec rules, not necessarily to identify the optimal configuration of IPsec options. Therefore, the term ‘significantly degrading’ must be qualified. For the purpose of this thesis, this simply means that any protection afforded by IPsec must be available after all streamlining decisions are made. Setting the choice of encryption algorithm to a default that is not recognized as the strongest encryption

algorithm available does not constitute a significant degradation in security. Encryption is still provided and the specific algorithm chosen can be changed later.

A good frame of reference for approaching this problem can be found in [32]. The authors' main purpose is to provide a cryptographic evaluation of IPsec, but they first address the protocol suite's complexity. As described above, the main parts of IPsec are evaluated regarding the complexity they bring, and the authors provide suggestions on how the protocol suite could potentially be simplified. It is important to note that the authors present only their feelings toward IPsec as a security protocol, and it is wise to consider both sides of any argument. Stephen Kent is the author of several IPsec-related RFCs, and a version of [32] that includes rebuttal comments by Kent is provided in [33].

Once again, this thesis is not arguing for or against any of these suggestions, nor does it support any of them as the optimal IPsec configuration. The purpose of this section is to evaluate the impact these decisions would have on IPsec administration and visualization. A full-featured version of the visual IPsec management application presented in this thesis should be able to support any IPsec implementation to some degree.

The following section evaluates several different areas of IPsec and identifies each decision made to simplify the implementation. The results of these decisions as they apply to the complexity of both management and the visualization of IPsec rules are presented in Section 4.1.

3.1.2.1 *Operating Modes*. IPsec must include Tunnel Mode to handle communication where the endpoints of the protected tunnel are not necessarily the endpoints of the communication channel. Recall from Chapter II that Tunnel Mode creates an inner and outer IP header. However, Tunnel Mode can be used for end-to-end communication by having these headers be the same. This seems like the ideal choice, but using Tunnel Mode for end-to-end communication introduces additional overhead with the additional header information. The authors of [32] suggest eliminating Transport Mode and using a header compression scheme to reduce the increased overhead of Tunnel Mode. In practice it might be prudent to implement this only after a compression scheme has been tested and accepted as a standard component of the IPsec protocol suite rather than implementing any random scheme available. However, for the purpose of illustrating the effects of simplifying the IPsec implementation, this choice works well. An alternate solution might be to have the system automatically use Transport Mode for communications where the source and destination addresses are within a specified range (typically the local area network) and Tunnel Mode for communications where either the source or destination address falls outside that range. This would achieve a similar effect regarding management simplification but would certainly be more complicated to implement.

3.1.2.2 *Protocols*. Both AH and ESP provide authentication, and ESP can additionally provide encryption. It is logical to question whether AH is required at all, especially in the context of simplifying IPsec. AH was maintained because it provided increased security in Transport Mode because it authenticated parts of the IP

header [24]. Following the recommendation to eliminate Transport Mode, eliminating AH makes sense since ESP authentication in Tunnel Mode is considered to be equally secure. When the IPsec RFCs were revised in 2005, AH was changed from a mandatory part of any IPsec implementation to an optional one further supporting the idea that ESP is sufficient in most cases.

3.1.2.3 *Encryption.* The authors of [32] stress that encryption without authentication is useless because, even though the data is encrypted, it is susceptible to certain types of attacks. They recommend that ESP always enforce authentication and make adding encryption optional. This is the preferred solution in conjunction with the two previous decisions. It further simplifies administration and enables a more straightforward visualization as will be illustrated later. If there are situations where encryption without authentication is desired or required, an actual implementation could allow an administrator to alter the settings for the specific instance.

3.1.2.4 *Algorithms.* IPsec is designed to support various cryptographic algorithms for both authentication and encryption. Several algorithms are designated in the RFCs as 'must' or 'should' be supported, and additional algorithms can be incorporated into an implementation. It is not practical to limit the algorithms available for either service for the sake of streamlining the interface. However, using the assumption that the majority of communication within a given local area network can be afforded the same level of protection, setting the authentication and encryption algorithms to a default simplifies the management overhead and allows the administrator to only need to select settings when something other than the typical case is required.

There are other areas of IPsec that could be evaluated but that have less of an effect on both the management and visualization of IPsec rules. Perhaps most notably is key exchange. It is sufficient here to say that key exchange should be handled in a standard default manner using IKE for key exchange. This eliminates the need for an administrator to make decisions on how to configure key exchanges and avoids the problems associated with manual keying. As always, an implementation could include functionality allowing an administrator to change these settings to handle specific cases.

3.2. Visualizing IPsec Rules

Visualization is a proven approach to dealing with complex datasets. Managing IPsec rules becomes increasingly complex as the number of nodes and the number of rules on the network increase. Given the success of visualization in the network management arena, visualization seems to be the logical, if not ideal, approach to managing IPsec rules.

3.2.1. Problem Definition

The problem is there is no efficient, intuitive method for managing complex sets of IPsec rules. To make IPsec deployment on production networks viable, administrators require management approaches that enhance situation awareness and simplify management functions.

3.2.1.1 *Goals and Hypotheses*. The main goal is to develop a visual representation of IPsec rules deployed on a network that increases situation awareness and eases management for an administrator. Another goal is for the visualization

developed to be intuitive to a network administrator as opposed to simply being a representation of IPsec rules that has no context on its own. The hypothesis is that visualization can improve both situation awareness and management of IPsec similarly to how it has been applied to network management.

3.2.2. Approach

Visualization is both art and science. Developing an effective visualization is not an endeavor that lends itself easily to many straightforward experimentation techniques. There may be certain aspects where it can be shown that a given technique is not suited to a given problem, but it is more difficult to know if a chosen approach is an effective one. Applying several visualization techniques to a dataset of IPsec rules produces various visualizations that can be evaluated and compared regarding their applicability to the problem at hand. Experience, background research, and correspondence with network and visualization professionals provide valuable insight into the effectiveness of the final visualization design decisions presented here.

3.2.2.1 *Visualizations*. The various visualization techniques described in Section 2.1.1 represent only a sampling of the possibilities for visualizing data. However, they provide a solid foundation for exploring the effectiveness of visualizing IPsec rules in different ways. This thesis presents visualizations of IPsec rules developed using the following techniques:

- Parallel Coordinate Graphs: this multivariate display shows IPsec connections between source and destination addresses showing the source and destination ports and the protocols for each rule

- Treemaps: explores the space-filling approach to represent IPsec rule information for all nodes in a confined space
- Glyphs: this approach encodes IPsec rule information onto the existing network map typical to a network management tool to present IPsec information to administrator in an immediately familiar environment
- Radial Link Graphs: visualizes IPsec connections as logical connections between nodes

Each of the visualizations produced must be evaluated to determine its effectiveness in managing IPsec rules on a network. The evaluation criteria described below were derived from the stated goals of the visualization and key concepts for creating a security visualization system found in [4].

- Does the visualization represent all data dimensions?
- Are data representations clear and easily distinguishable?
- Is the visualization intuitive to network administrators?
- Is the visualization scalable?
- Can specific information be found quickly?
- Is the visualization approach suitable for managing data on a production network?

This thesis provides only an initial evaluation if the visualizations against these criteria. The results would require validation through a human user study where people of various backgrounds and experience levels could explore the visualizations and provide feedback. Time constraints prevented such a study from being accomplished for

inclusion in this research effort. The initial evaluation of each visualization is presented in Chapter IV.

In order to be able to effectively compare the different visualizations produced, a simple network was designed and IPsec rules were developed for each node to provide a single dataset for all visualizations. The network is described in the next section.

3.2.2.2 *Network Design.* A simple local area network was designed to apply a baseline set of IPsec rules to. Having an established set of IPsec rules allows multiple visualization techniques to be explored producing different visualizations of the same data that can withstand direct comparison.

The network was designed to be simple yet realistic. Some characteristics of the network design were ignored (such as the types of routers or switches used) since they did not affect the IPsec rules on the network. The IPsec rules applied follow the design decisions suggested in Section 3.1, so all rules are using the ESP protocol in Tunnel Mode. The network consists of 32 nodes. They are listed below with a summary of the IPsec rules applied to each. The specific rules for each node are provided in Appendix A.

- Domain Controller: authenticates with all systems in the domain (in this example, this excludes the Web server and External DNS server) on any port using any protocol
- DNS server: authenticates with all systems except the Web server on port 53 over UDP and TCP
- Mail server: authenticates with all systems in the domain except for the file server on any port using any protocol

- File server: authenticates with all workstations and encrypts traffic with some workstations on any port using any protocol
- External DNS server: authenticates with the Web server and internal DNS server on port 53 over UDP and TCP
- Web server: authenticates with the workstations on port 80 over HTTP and port 443 using HTTPS
- Workstations (26): authenticate and encrypt traffic between workstations on various ports over multiple protocols

3.2.2.3 *Modeling*. The Prefuse Information Visualization Toolkit [34] is used to model various visualizations. The toolkit is written in Java and is designed for data modeling, visualization, and interaction. It enables the visualizations developed here to be modified quickly and provides a clear idea of what interacting with the visualizations would be like versus having only static images. The toolkit also has additional tools that can be applied to the visualizations created providing additional views, insights, and potential solutions to issues or directions to pursue.

3.3. Summary

This chapter identified the areas of focus for this thesis: streamlining the IPsec implementation and visualizing IPsec rules to ease the burden of managing IPsec when deployed on a local area network. It defined the problem in each area and outlined the approaches used to address each issue. The results of these approaches are presented in Chapter IV.

IV. Results

This chapter explores the results of the approaches outlined in Chapter III. Section 4.1 identifies the impact the decisions for streamlining the IPsec implementation have on both IPsec management and on the visualizations that might be used to represent IPsec rules. Section 4.2 presents the various visualizations of the IPsec rules listed in Appendix A, which were provided for the network described in Section 3.2.2.2. Section 4.3 explores how the visualization of IPsec rules could be employed in a network management tool to enhance an administrator's situation awareness and ease the management of IPsec rules deployed throughout a network.

4.1. Simplifying IPsec

The implementation decisions outlined in Section 3.1.2 have a major impact on the complexity of both managing and visualizing IPsec rules. The idea that reducing the complexity of IPsec would in turn reduce the complexity of IPsec management has been discussed in the literature throughout IPsec's evolution [24] [26] [32] [33], but it is usually stated as a simple truth that people agree on in general. In this section, the impact of the streamlining decisions is explicitly quantified. Additionally, since the focus of this research is employing visualization to manage the dataset, the effects of the decisions on each of the visualization approaches is also examined. Although the design decisions here are specific to IPsec, this evaluation illustrates the importance of considering the impact that design decisions have on various aspects of a system such as how data might be managed or represented.

4.1.1. Simplifying IPsec Configuration

An administrator has to make several decisions or specify many configuration settings for every IPsec rule that needs to be configured. Assuming there are several nodes throughout a given network that require similar or even identical levels of protection, even having options set to defaults and only requiring actions when settings need to change would ease administration. Streamlining the implementation and setting algorithm options, which should not be limited to a single choice, to defaults simplifies configuration tasks for the administrator.

To create an IPsec rule between two nodes, an administrator must specify five pieces of information to identify traffic that requires IPsec protection: source and destination addresses, source and destination ports, and a communication protocol. This information is necessary regardless of whether the implementation has been streamlined or if all IPsec components are available.

However, without the streamlining decisions outlined in Section 3.1.2, the following additional decisions must be made:

- Which operating mode will be used: transport mode or tunnel mode?
- Will the traffic be authenticated? If yes,
    - Which protocol will provide authentication: the AH or ESP protocol?
    - Which authentication algorithm will be used?
- Will the traffic be encrypted? If yes,
    - Which encryption algorithm will be used?

Six decisions must be made for every IPsec rule in addition to the initial five pieces of data an administrator must provide for a total of 11 data points. While this may not seem significant on its own, consider the data used to generate the visualizations in this thesis. IPsec rules were generated for 32 network nodes. While these rules were intended to be similar to what might be found on a production network, actual application and security requirements could result in a much larger set of rules than what was used. Even so, a network with only 32 nodes generated 542 distinct IPsec rules (see Appendix A). If every rule was configured manually, there would be over 2,000 rules due to the fact that there are duplicate rules for inbound and outbound traffic on each node and each rule pair needs to exist on both communication endpoints. Even on a 32-node network with just a few hundred rules to configure, an additional six decisions to make is unacceptable if it could be easily avoided. Consider the overhead these additional decisions would cause when configuring a network with several hundred nodes.

Recall from Section 3.1.2 that the suggested changes to IPsec included eliminating transport mode, eliminating the Authentication Header protocol, and always providing authentication. Implementing these three changes reduces the number of additional decisions to be made by 50%, leaving only the following decisions for the administrator to make:

- Which authentication algorithm will be used?
- Will the traffic be encrypted? If yes,
    - Which encryption algorithm will be used?

Section 3.1.2 also suggests that the authentication and encryption algorithms be set to a default so an administrator would only need to take action under special circumstances. Following the assumption that most nodes on a given network would require the same level of protection regarding the algorithms used, setting the algorithms to default values potentially eliminates the need to specify authentication and encryption algorithms for the majority or rules established. This would mean that an administrator would now only need to specify one additional piece of information for most rules:

- Will the traffic be encrypted?

This reduces the additional workload by approximately 83% and cuts the total number of data points that need to be provided for each rule from 11 to 6 for an overall reduction of 45%. Using the network described in Section 3.2.2.2, which consists of a total of 2,176 IPsec rules (remember there are four versions of each of the 542 distinct IPsec rules) using all default settings, this equates to a maximum of 2,168 extra actions (if encryption was selected for all rules) versus a maximum of 13,008 actions if all six potential configuration options had to be specified.

4.1.2. IPsec Visualizations

More significant to this research are the effects of the implementation details on the various visualizations of IPsec rules presented here. This section provides examples of visualization techniques using a small dataset as it would look showing all IPsec options followed by a representation of a similar dataset employing the design decisions in Section 3.1.2. The datasets provide the same number of IPsec rules configured

between the same number of nodes. The only differences between the datasets are the first uses both protocols for authentication, both operating modes, and a mix of protocols when encryption and authentication are used.

4.1.2.1 In order to fairly compare visualizations of standard and simplified IPsec rules, the simplified dataset must provide the same protections rule for rule as the standard dataset while employing the streamlining decisions described in Section 3.1.2. Table 5 provides a simple set of rules between 8 network nodes. Table 6 implements each rule in Table 5 under the restrictions imposed by the decisions in Section 3.1.2. Only distinct outbound rules are represented to save space. Mirrored inbound rules and reciprocal rules on destination nodes are not included.

Table 5: Sample IPsec rules for 9-node network with no streamlining

Rule	Src Add	Dest Add	Src Port	Dest Port	Protocol	Action	Auth Prot	Enc Alg	Auth Alg	Mode
1	Node 1	Node 2	ANY	ANY	ANY	IPsec	AH	NULL	HMAC-SHA-1	Tunnel
2	Node 1	Node 3	ANY	ANY	ANY	IPsec	AH	NULL	HMAC-SHA-1	Tunnel
3	Node 1	Node 4	ANY	ANY	ANY	IPsec	AH	TripleDES-CBC	HMAC-SHA-1	Tunnel
4	Node 1	Node 7	ANY	ANY	ANY	IPsec	ESP	TripleDES-CBC	HMAC-SHA-1	Tunnel
5	Node 2	Node 3	ANY	ANY	ANY	IPsec	ESP	NULL	HMAC-SHA-1	Transport
6	Node 2	Node 5	ANY	ANY	ANY	IPsec	AH	NULL	HMAC-SHA-1	Transport
7	Node 2	Node 6	ANY	ANY	ANY	IPsec	AH	NULL	HMAC-SHA-1	Transport
8	Node 4	Node 5	ANY	ANY	ANY	IPsec	AH	NULL	HMAC-SHA-1	Tunnel
9	Node 4	Node 6	ANY	ANY	ANY	IPsec	AH	NULL	HMAC-SHA-1	Tunnel
10	Node 4	Node 7	ANY	ANY	ANY	IPsec	NULL	TripleDES-CBC	HMAC-SHA-1	Tunnel
11	Node 5	Node 6	ANY	ANY	ANY	IPsec	ESP	NULL	HMAC-SHA-1	Transport
12	Node 5	Node 8	ANY	ANY	ANY	IPsec	AH	NULL	HMAC-SHA-1	Transport
13	Node 5	Node 9	ANY	ANY	ANY	IPsec	AH	NULL	HMAC-SHA-1	Transport
14	Node 7	Node 8	ANY	ANY	ANY	IPsec	AH	NULL	HMAC-SHA-1	Tunnel
15	Node 7	Node 9	ANY	ANY	ANY	IPsec	AH	NULL	HMAC-SHA-1	Tunnel
16	Node 8	Node 9	ANY	ANY	ANY	IPsec	ESP	NULL	HMAC-SHA-1	Transport
17	Node 8	Node 2	ANY	ANY	ANY	IPsec	AH	NULL	HMAC-SHA-1	Transport
18	Node 8	Node 3	ANY	ANY	ANY	IPsec	AH	NULL	HMAC-SHA-1	Transport

Table 6: IPsec rules from Table 5 with streamlining decisions incorporated

Rule	Src Add	Dest Add	Src Port	Dest Port	Protocol	Action	Enc Alg	Auth Alg
1	Node 1	Node 2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
2	Node 1	Node 3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
3	Node 1	Node 4	ANY	ANY	ANY	IPsec	TripleDES-CBC	HMAC-SHA-1
4	Node 1	Node 7	ANY	ANY	ANY	IPsec	TripleDES-CBC	HMAC-SHA-1
5	Node 2	Node 3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
6	Node 2	Node 5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
7	Node 2	Node 6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
8	Node 4	Node 5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
9	Node 4	Node 6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
10	Node 4	Node 7	ANY	ANY	ANY	IPsec	TripleDES-CBC	HMAC-SHA-1
11	Node 5	Node 6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
12	Node 5	Node 8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
13	Node 5	Node 9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
14	Node 7	Node 8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
15	Node 7	Node 9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
16	Node 8	Node 9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
17	Node 8	Node 2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
18	Node 8	Node 3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1

As you can see, Table 5 requires a column for specifying the authentication protocol used, if any, and a column for specifying the operating mode used. In Table 6, authentication is always provided using the Encapsulating Security Payload (ESP) and Tunnel Mode is always used eliminating the need for those columns. Otherwise, each rule entry in Table 6 corresponds to the same entry in Table 5 regarding the IPsec protection provided to the specified traffic. Take note that these rules are strictly for illustrative purposes only and are not meant to imply an actual working or suggested network configuration.

To explore the effects the streamlining decisions have on potential data representations, the rules in Table 5 and Table 6 are rendered using two visualization approaches: treemaps and radial link-node graphs. Both approaches were introduced in Section 2.1.1, and more detailed examples are evaluated in Section 4.2 using the data

provided in Appendix A. Each is only briefly discussed here to illustrate the effects that the differences in the datasets have on the data representations.

4.1.2.2 *Treemaps.* Recall from Section 2.1.1.3 that treemaps are space-filling approaches using nested rectangles alternating horizontally and vertically as new data points are represented. Figure 26 is a treemap representation of Table 5. Each node was given the same amount of space for this example. First, the space for each node was divided horizontally to represent each authentication protocol used for rules on a given node. Each subdivision was then further divided vertically to account for each rule using the identified authentication protocol. Encryption is represented by a red highlight, and shading is used to identify where Tunnel mode is used as opposed to Transport mode.

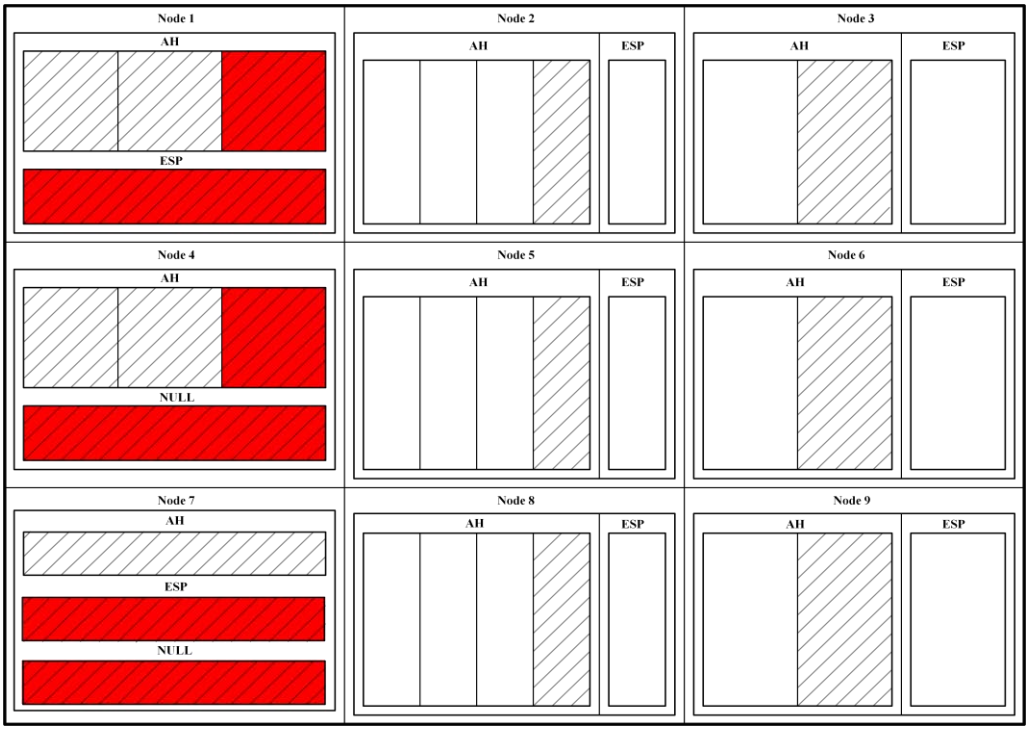

Figure 26: Treemap representation of Table 5 rules with no streamlining

Each node needs to present up to four data dimensions: authentication protocols, rules within each authentication protocol, whether or not encryption is used, and which operating mode is used. Keep in mind that the rules in Table 5 use only TCP for simplicity, but the communication protocol would also need to be represented. This means each authentication protocol subdivision would first be divided vertically to represent the various communication protocols (ANY, TCP, UDP, etc.) and then further subdivided horizontally to represent each individual rule using the specified protocol.

Figure 27 is a treemap representation of Table 6. Each node is again given equal space in the display. However, now only two data dimensions need to be represented. The space for each node is divided vertically to create a space for each distinct rule. Again, color is used to identify encryption.

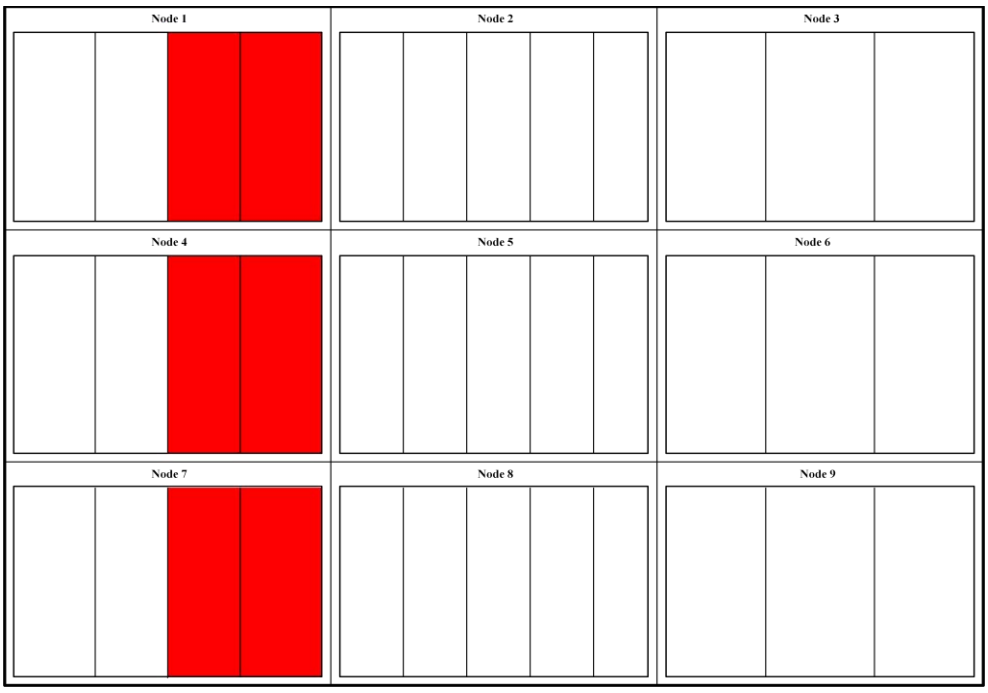

Figure 27: Treemap representation of Table 6 rules using streamlining decisions

The streamlining decisions result in fewer alternations between horizontal and vertical spaces and no need to distinguish between modes. As with Figure 26, even the streamlined visualization in Figure 27 would need to represent the communications protocols being used. This translates to Figure 27 realistically being expected to look like Figure 26 without the shading to represent the mode being used. This is still desired over adding an additional layer of subdivisions to Figure 26.

4.1.2.3 *Radial Graphs*. Radial graphs (Section 2.1.1.1) are good at representing relationships between network nodes. For this thesis, those relationships are the logical communication paths between nodes using IPsec. Figure 28 is a radial graph representation of the rules in Table 5. The data dimensions for Figure 28 break down as follows:

- Black lines indicate that only authentication is provided
- Red lines indicate that encryption is used
- Thick lines represent Tunnel mode
- Thin lines represent Transport mode
- Solid lines indicate authentication is provided using AH
- Dashed lines indicate authentication is provided using ESP
- Dotted lines indicate no authentication is provided

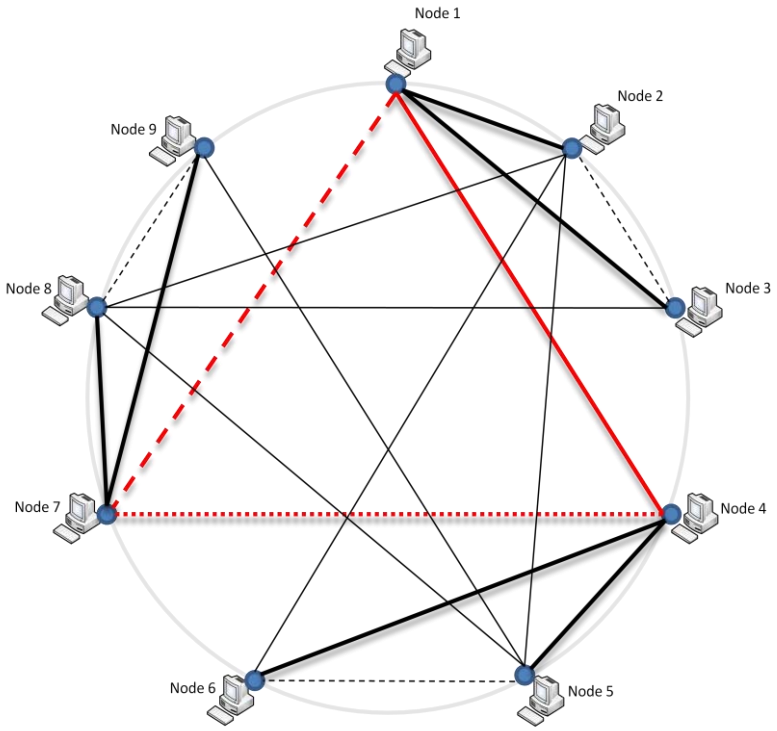

Figure 28: Radial graph representation of Table 5 rules with no streamlining

Figure 28 shows only one of many possible ways to represent the data in question. However, even this simple example raises some concerns. For example, line thickness is used to distinguish between Transport mode and Tunnel mode. The difference is fairly obvious when looking at the black lines. However, if only one mode was represented like with the red lines that all use Tunnel mode, could a viewer know immediately which is being represented? Also, there might be some confusion when judging the thickness of the dotted line between Node 4 and Node 7. The dotted line may appear thinner to some due to the smaller segments being used even though the line thickness is set to the same value as the other red lines.

Figure 29 is a radial graph representation of the rules in Table 6. Since authentication is always provided using ESP and Tunnel mode is always used, there are only two data dimensions to represent:

- Black lines indicate that only authentication is provided
- Red lines indicate that encryption is used

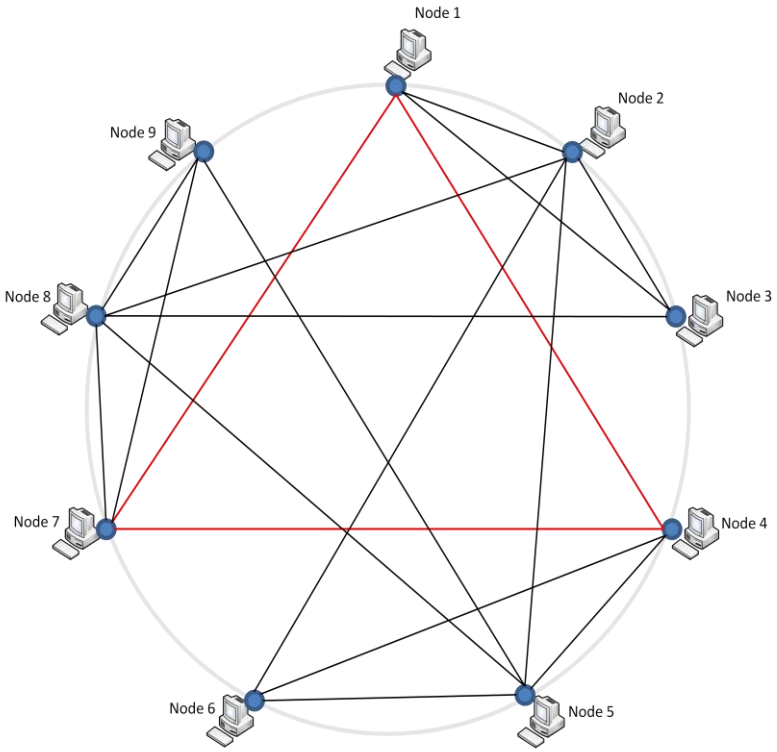

Figure 29: Radial graph representation of Table 6 rules using streamlining decisions

Clearly the streamlined IPsec implementation leads to a more straightforward, useable visualization. With only two data dimensions to represent, an administrator can look at any link and immediately know the one distinguishing feature about the IPsec rule in question: is encryption used or not? The only additional information about a rule would be what the source and destination ports are and which algorithms were being

used. Since there are a greater number of choices, it is not practical to try to represent this information immediately through the visualization. It is important to note that the visualization scheme could be used for the data in Table 5, but the additional data dimensions would need to be provided somewhere. It could be presented in conjunction with the port and algorithm information, but that would defeat the purpose of the visualization which is to present as much information as quickly and clearly as possible.

Considering the impact the streamlining decisions had on the simple dataset used in the examples presented in this section, their benefit when visualizing more complex, realistic scenarios should be easy to see. Look at some of the visualizations presented in Section 4.2, which are streamlined visualizations of 32 nodes versus nine and 542 distinct IPsec rules compared to 18, and try to picture them with the additional complexity described in this section. From a visualization perspective, simpler is definitely better.

4.2. Visualizing IPsec Rules

IPsec is complex, and managing IPsec on a network becomes increasingly complex as the number of rules increases. When developing the interface simulation, a 12-node network was initially used. An arbitrary set of rules was developed to closely mimic what might be found on an actual network with the same setup. This resulted in 66 distinct IPsec rules that needed to be represented. When the network design was expanded to 32 nodes by adding 20 workstation nodes with similar rules, the number of distinct rules jumped from 66 to 542. According to LTC Greg Conti, author of "Security Data Visualization", [4] and other visualization professionals, it does not appear that

visualization has ever been applied to IPsec for data representation or management simplification despite its complexity [36].

To evaluate visualization's effectiveness in representing and potentially managing IPsec data, the IPsec rules in Appendix A are represented using four different visualization techniques. The first approach, parallel coordinate graphs, and the second approach, treemaps, allow for a large amount of data to be presented simultaneously in a small space where the viewer can select specific pieces of data to view. The third approach examines the feasibility of encoding IPsec information onto an existing network management map using glyphs. Finally, the data is represented using radial link-node graphs that represent the IPsec rules as logical connections between network nodes.

4.2.1. Parallel Coordinate Graphs

Parallel coordinate graphs (Section 2.1.1.1) display multivariate data by plotting various data points across multiple axes. A large dataset can be represented in a relatively small space with an interface allowing the viewer to focus on specific pieces of data. This section explores representing IPsec data using a parallel coordinate graph and its potential as a management approach for that data.

4.2.1.1 *Visualization*. IPsec consists of several data points that are well suited for display on a parallel coordinate graph. Figure 30 presents a basic parallel coordinate graph for the IPsec data in Appendix A. The data points of source address, source port, protocol, destination port, and destination address are each represented by a separate axis. A line is drawn for each distinct IPsec rule starting at the source address and connecting all appropriate data points to the destination address. The connected

segments from source to destination address provide all five data points for a specific rule.

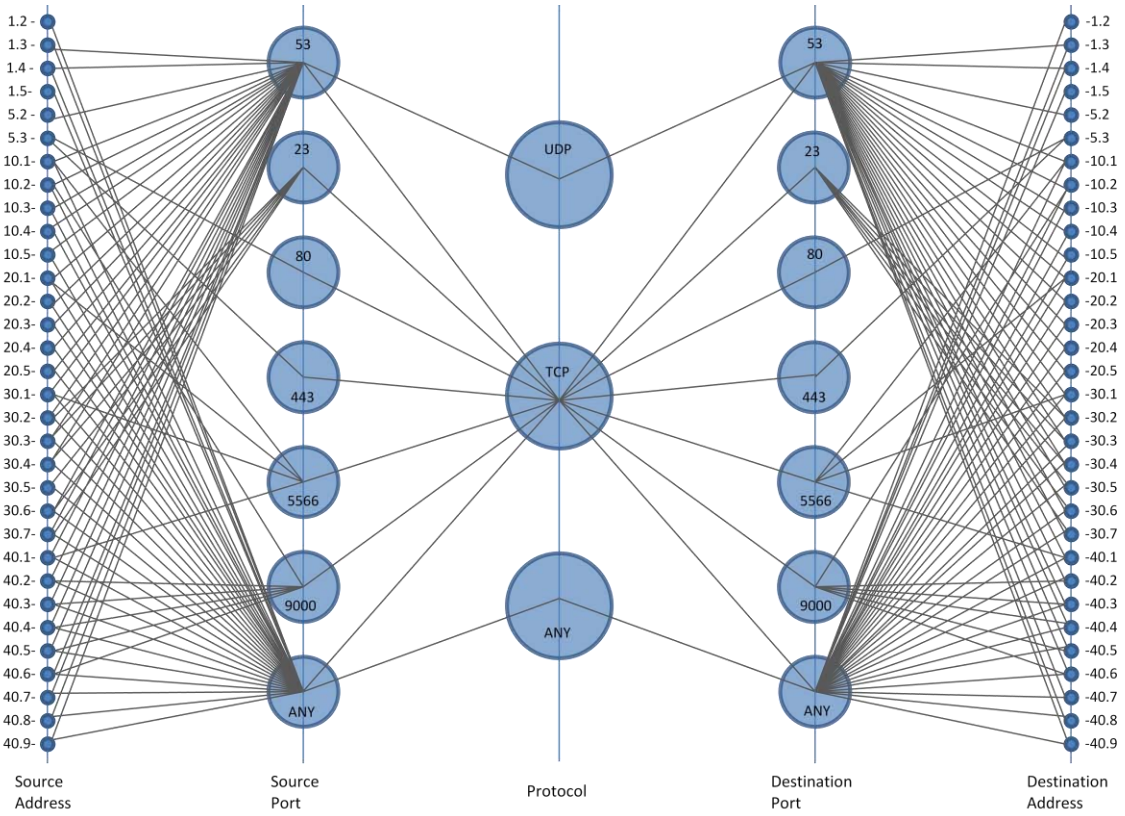

Figure 30: Parallel coordinate graph of IPsec rules in Appendix A

This parallel coordinate graph seems fairly easy to read considering that even in areas where there is a lot of congestion, any individual line segment can typically be followed from endpoint to endpoint with little effort. Unfortunately, there are two issues that limit the graph's usefulness as is. First, there is no way to tell for a single rule which individual line segments should be connected from axis to axis. Second, to keep the graph readable, duplicate lines were overwritten as opposed to being rendered multiple

times. For example, the internal DNS server (1.3) has two rules with Workstation 1 (10.1) across port 53: one for UDP and one for TCP. Rather than draw two lines from 1.3 to source port 53, only one line is rendered. As a result, it may be misinterpreted that 1.3 only has a single rule using source port 53 when in fact it has one for each network workstation. However, recall from Section 2.1.1.1 that this is to be expected. The usefulness of such a graph comes from interactivity that allows a viewer to select specific data points or groups of data to view. Working within the display, the interface could allow the viewer to select a specific data point and highlight all rules containing that data point. This is illustrated in Figure 31 where source address 1.3 has been selected.

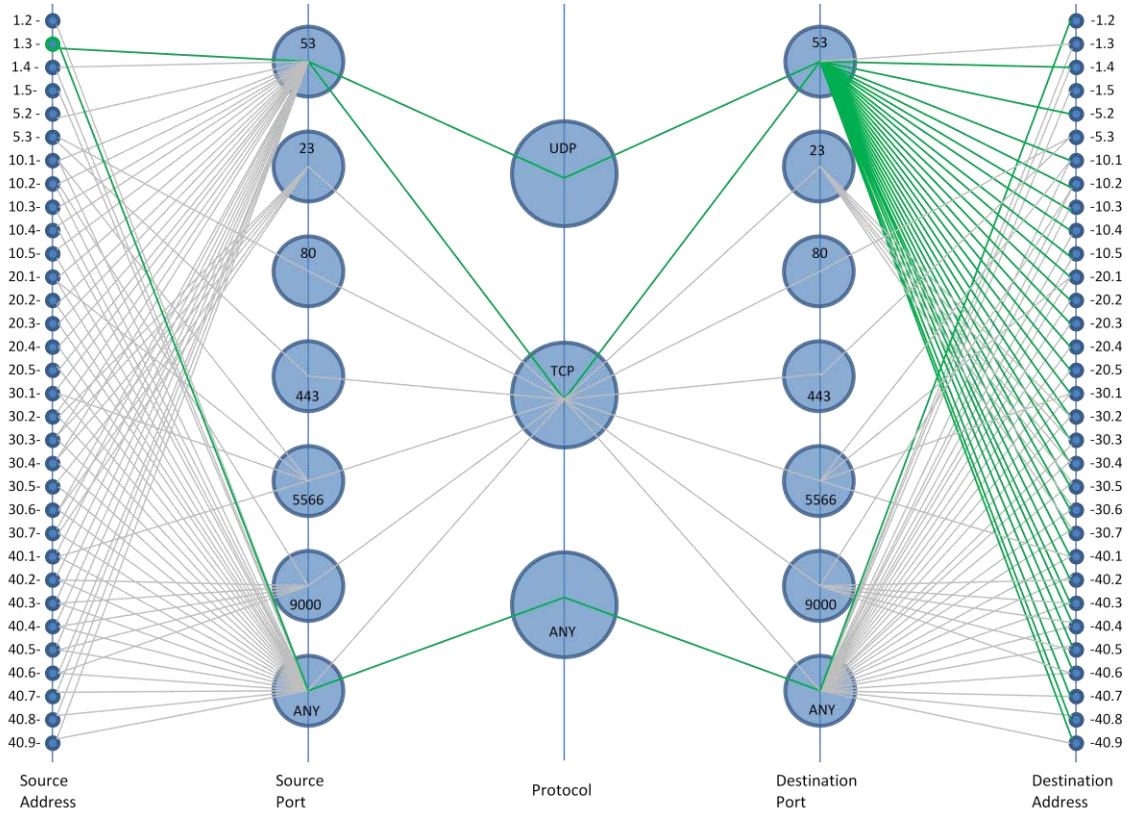

Figure 31: Parallel coordinate graph of IPsec rules with a specific source address selected and all associated rules highlighted

In Figure 31, the selected node is highlighted along with all line segments representing rules with 1.3 as the source address. To enhance clarity, the line segments not associated with 1.3 have faded, allowing the viewer to focus on the items of interest.

Through this approach, it is clear there is a rule covering traffic on any port using any protocol between source address 1.3 and destination address 1.2. There is still confusion though as to which segments leaving destination port 53 pertain to UDP versus TCP since both protocols feed into the same location. This presents an important lesson to be learned. Even if all of the data can be presented simultaneously, it may only be practical to manage limited subsets of the information at a time. To address this in Figure 31, the interface could allow subsequent selections of data points to further narrow the scope of information being viewed. For example, in Figure 31 the viewer might be able to select a source port followed by a protocol, etc. in addition to a source address to eliminate potential overlap.

The effectiveness of the display can be enhanced even further. When a specific data point is selected, the display could change using a simple animation that moves the highlighted rules and associated data points into focus while pushing the background noise out of the way completely. Figure 32 illustrates what Figure 31 might look like employing this technique. The information presented is the same, but the graph is much cleaner and all focus is now on the items of interest. Rather than condense the background noise to the bottom of the axes, the unused icons could simply fade from the graph to provide more display space.

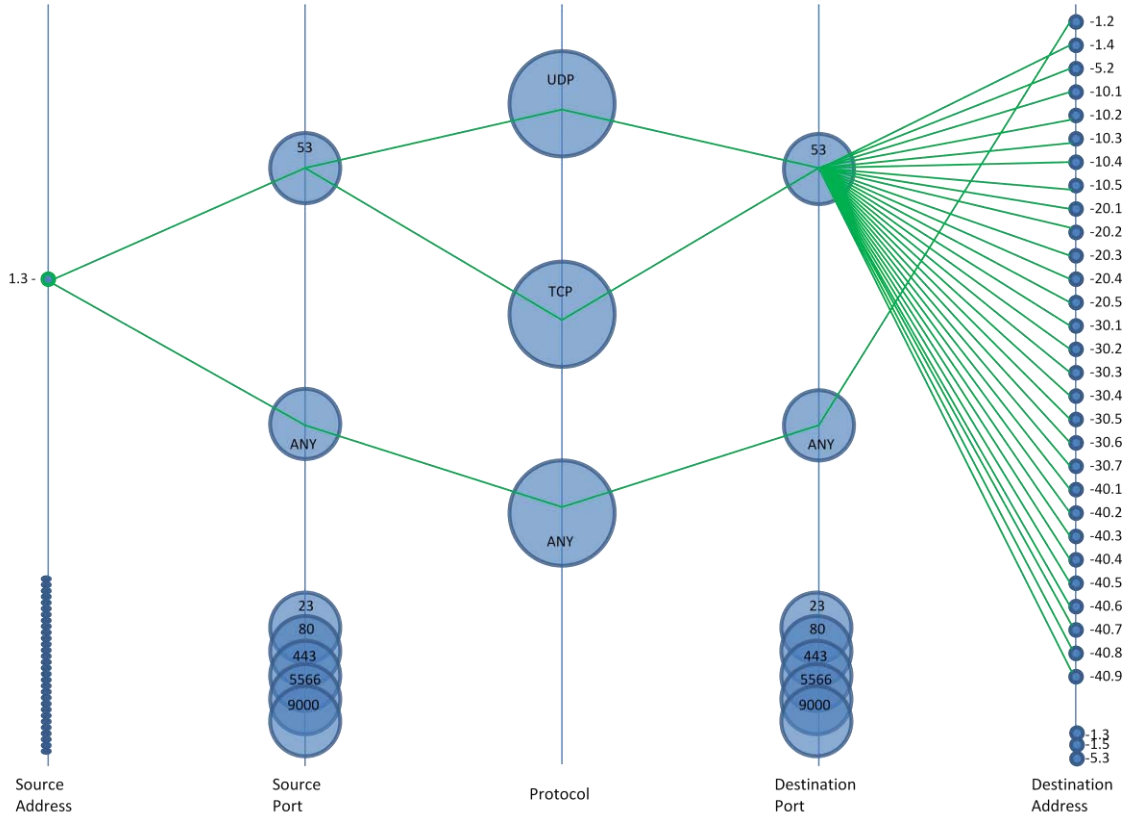

Figure 32: Result of animated parallel coordinate graph focused on data related to selected source address

Using the parallel coordinate graphs presented here, the five data points of source and destination address, source and destination port, and communication protocol can be presented simultaneously for a potentially limitless number of rules. There are two additional data points that need to be addressed; encryption and the algorithms used.

Within the graph, axes could be added to present values for the authentication and encryption algorithms. This thesis suggests that these algorithms be set to default values for the common case. This would mean that most rules would likely converge into the same algorithm entries on these additional axes. This does not seem to be a worthwhile tradeoff considering the additional clutter that would be added to the graph. For our

purposes, having the algorithm information show up in a window when a specific rule is highlighted or selected is more desirable. Whether or not a rule provides encryption versus authentication only is something an administrator might what to know at a glance. This can be addressed in different ways, but many methods would also involve additional clutter on the graph. For example, any ports and protocols that pass both encrypted and non-encrypted traffic could simply be duplicated on their respective axes. Another approach could provide a split display showing both types of traffic for selected data points. These approaches could be effective but require additional space or complexity in the display. To avoid this, the information would need to be encoded into the existing components of the display. One possible way to do this is with color.

Figure 33 illustrates how color can be used to differentiate between authentication only, rules that provide encryption, and line segments that represent both types of rules at the same time. In this example, a web server provides IPsec rules for data on port 80 and port 443 to all workstations. For traffic over port 443, IPsec provides additional encryption.

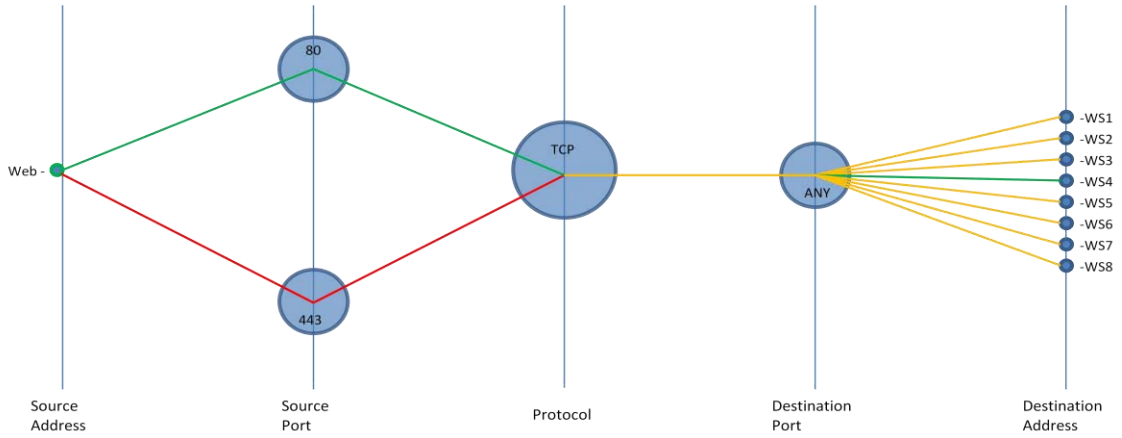

Figure 33: Using color to multiple data values onto single line segments: green represents authentication only, red represents encryption, and yellow represents both types of rules

Traffic from the web server over port 80 using TCP to any port on the workstations is colored green. Traffic from the web server over port 443 using TCP to any port on the workstations is colored red. To avoid rendering duplicate lines, the overlapping line segments are represented by a third color: yellow. An administrator can immediately see that WS4 is missing a rule for traffic over port 443. This may be intentional or may prompt the administrator to take action.

Using the techniques described in this section for focusing on specific information, Figure 34 shows the change in the graph in Figure 33 if the viewer were to select source port 443 from the display. In this case, the animation would eliminate information pertaining to port 80 allowing the view to focus on the singular rule type that remains. Note that WS4, the node in Figure 33 that had no rule for port 443, has been removed from the graph.

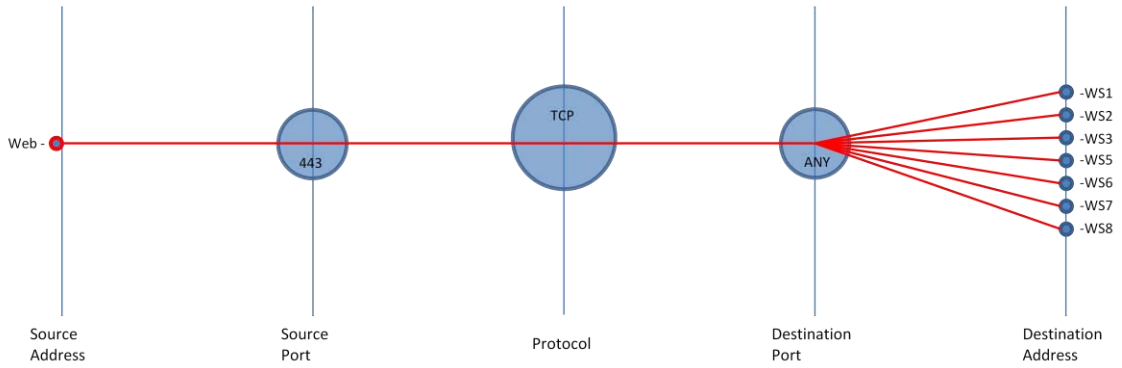

Figure 34: Sample state of display after narrowing focus to single rule type

4.2.1.2 *Evaluation.* Each visualization must be evaluated against the criteria laid out in Section 3.2.2.1. The parallel coordinate graph does well representing all of the data dimensions. The only data dimension not represented in the examples in this section is the algorithms used by each rule. While this data could be included on a separate axis, it is more efficient to present this information in a pop up window for a selected rule. The distinctions between different data representations are clear. This approach does not present the data in a way common to typical network management tools. However, the parallel coordinate graph is essentially just a line graph, which most people are familiar with, and the data is presented clearly and logically enough that an administrator could be expected to decipher the data fairly easily. The visualization is definitely scalable. There is no obvious limit to the number of rules a parallel coordinate graph could represent. Even with methods of isolating data like the ones suggested here, enough data could be provided to prevent all items of interest from fitting within a single display. Additional methods such as scrolling the graph would need to be used as well. It is simple to focus on specific information through various potential interface options.

Considering all these factors, it seems feasible to manage data on a production network using this approach. The only major drawback seems to be the display's lack of the feel of a network management application.

4.2.2. Treemaps

Treemaps (Section 2.1.1.3) are a space-filling approach to representing multivariate data using rectangular subdivisions of space alternating horizontally and vertically. This section explores the results of visualizing IPsec data using treemaps.

4.2.2.1 *Visualization*. Like the parallel coordinate graph, treemaps can encode a large volume of data within a relatively small space. However, treemaps are relatively new and likely to be unfamiliar to an administrator. As a result, it is especially important to display the right information as clearly as possible and to develop interface features that help the viewer navigate through the data.

Figure 35 is a treemap representation of the 542 distinct IPsec rules provided in Appendix A. For this example, size was not used to indicate a data point such as number of rules or percentage of rules providing encryption. For clarity, equal space was given across the top for the six network servers, and the remaining space was divided as equally as possible for the 26 workstations. The space for each node was then divided horizontally for each protocol covered by IPsec rules. The space within each specific protocol area was divided equally to provide a representation of each individual rule. Finally, any rules that provide encryption are identified by red highlight.

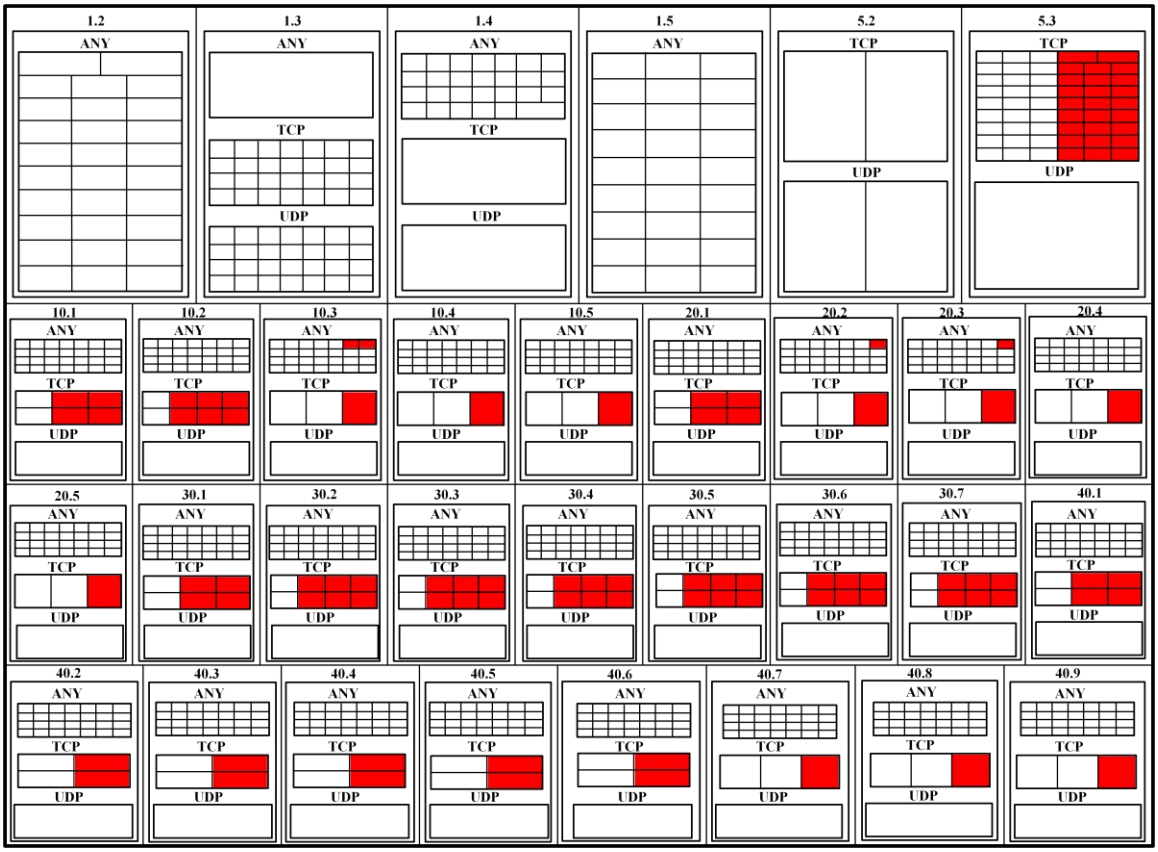

Figure 35: Treemap showing IPsec rules broken down by protocol for each network node

There is not enough available space to encode information such as destination address, ports, or algorithms directly on the display. The interface, however, would help mitigate this. Moving the cursor over or selecting a specific rule could highlight the rule and corresponding destination rule as well as bringing up a display window with additional information. This is illustrated in Figure 36. Other areas of the interface could allow for additional methods of selecting or displaying data as well. For example, a single node might be selected resulting in the node and all associated rules on other nodes

being highlighted. As with the parallel coordinate graph, animation could be used to redraw the display focusing only on the items of interest.

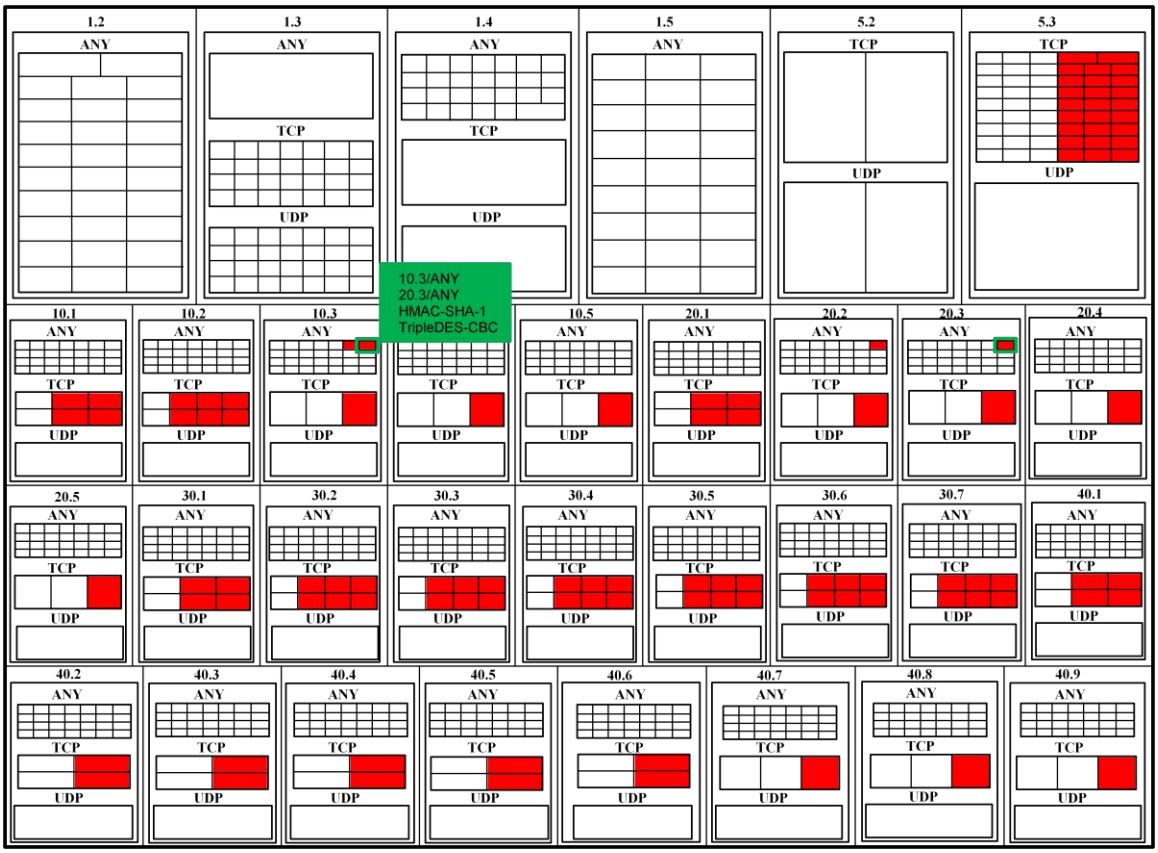

Figure 36: Viewing specific IPsec rule information using a treemap

4.2.2.2 *Evaluation.* The treemap approach enables quick presentation of massive amounts of data. While this approach is capable of representing the required data dimensions, less data can be explicitly show as the number of distinct entries in the map increases. It is also fairly simple to distinguish between data items, but this can also become more difficult as the number of entries increases. Unfortunately, this approach is not likely to be immediately familiar to network administrator and could require a steep

learning curve compared to other approaches. Treemaps scale well, often presenting thousands of pieces of data in a relatively small space, but they rely on an effective interface to navigate through the data. Information can typically be found quickly even with a very large dataset, though this is also directly related to the features of the interface. While possible interface functionality gives this approach potential, the overall foreignness of the display keeps this technique from being a more viable option.

4.2.3. Glyphs

Glyphs (Section 2.1.1.2) are graphical representations of multivariate data. Key features of glyphs such as size, shape, or color are used to uniquely represent different data points. This section explores the viability of encoding IPsec data onto a physical network map using glyphs.

4.2.3.1 *Visualization*. Figure 37 represents a potential physical view of the network described in Section 3.2.2.2 that might be used to manage the nodes on the network. As discussed in Chapter II, various icons are used to distinguish between different types of network nodes, and each node is labeled with its IP address. Nodes might change color to indicate changes in status, and an administrator could likely mouse over or right click a node to obtain additional information or take an administrative action. We can assume we can right click on a node and select an option to manage IPsec rules as suggested in Figure 37.

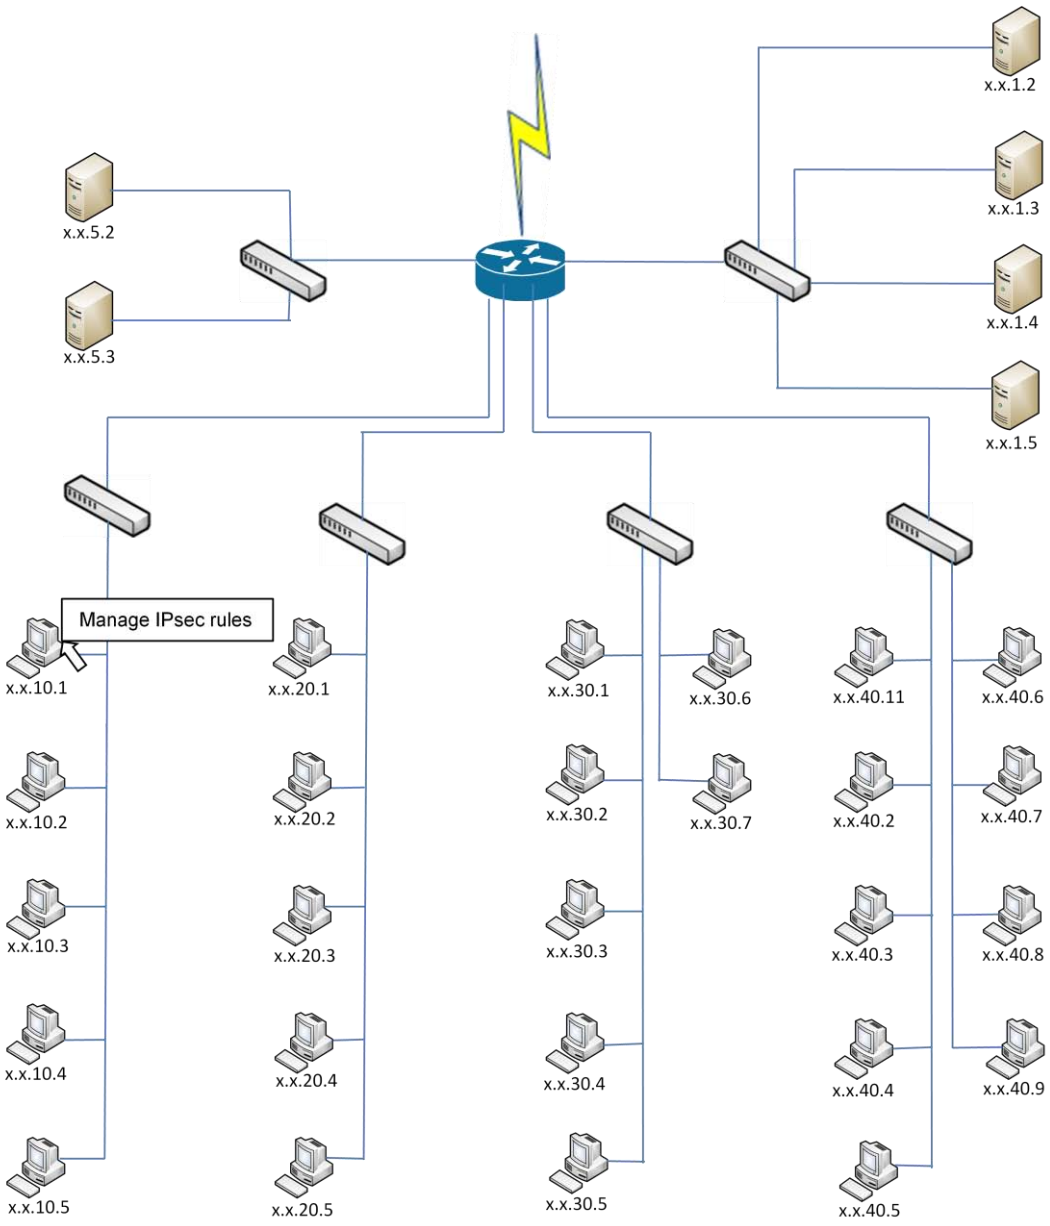

Figure 37: Physical map of network described in Section 3.2.2.2

A major challenge in using glyphs to represent data is selecting an appropriate glyph to represent the desired data dimensions. For this example, we use a circular glyph divided like a pie chart with each slice representing an individual IPsec rule on the selected node. Red highlights indicate rules providing encryption. Figure 38 shows a

possible result of displaying IPsec data for each node on the network using this approach. With so little space available, it is not practical to encode additional information within each individual slice, so moving the cursor over a slice could present additional information pertaining to the specific rule in a window as illustrated in Figure 38.

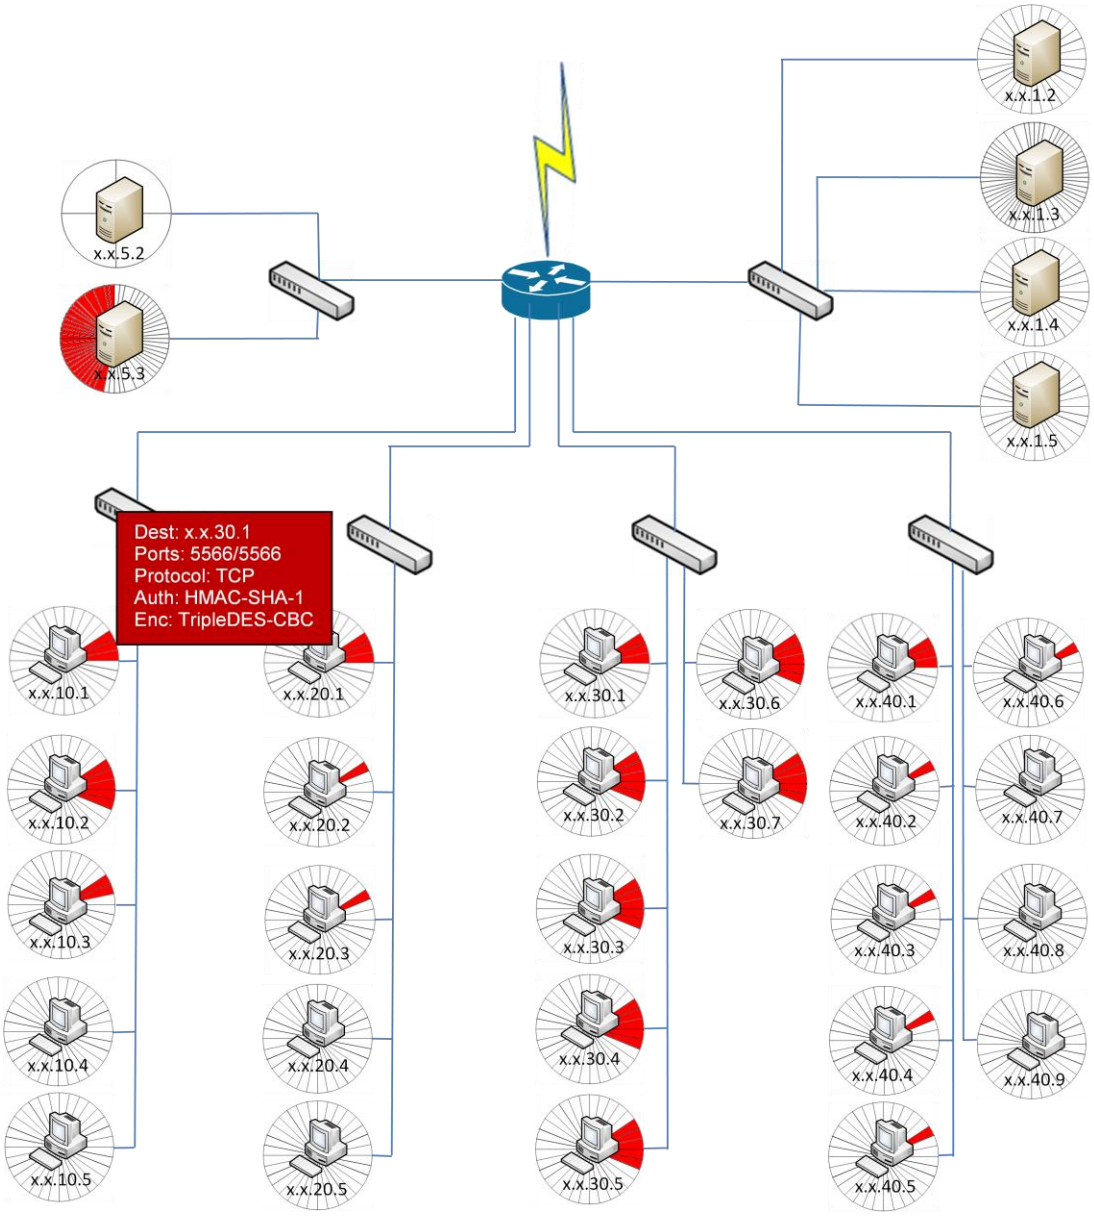

Figure 38: IPsec information encoded onto physical map using glyphs

There are several issues that immediately become apparent with this approach. Space is obviously an issue. The example above depicts a small sample network. If the network consisted of several hundred nodes, not only would each icon require additional display space to accommodate the glyph, but it would also become increasingly difficult to distinguish between individual slices within the glyphs as the number of rules on a single node increased. The is also little aggregate information that can be quickly discerned using this approach other than an idea of the number of encrypted versus non-encrypted rules in use. The inability to encode rule information within the individual slices that represent the rules further limits the amount of data that can be quickly presented to the viewer. Most data dimensions are restricted to being displayed in the pop-up window when the cursor is placed over an area of the glyph.

Like the previous methods, administrative functionality can address some of this approach's shortcomings. For example, selecting an individual slice could highlight the corresponding slice on the destination node, or selecting a node could highlight all of the destination nodes that share an IPsec rule with the selection. However, unlike the previous approaches, it would be more difficult to refocus the display on the items of interest and still maintain the advantages provided by having the physical layout.

Another approach would be to use glyphs on a physical network map to present IPsec information only on one node at a time. Figure 39 shows the IPsec rules assigned to a selected node. Again, moving the cursor over a slice presents the additional rule information pertaining to that specific rule. Additionally, selecting a rule from the glyph has highlighted the destination node on the map. The destination node might not always be located immediately with the display window, but this could be more easily mitigated

for a single source/destination pair by zooming out or perhaps having the path between source and destination highlighted.

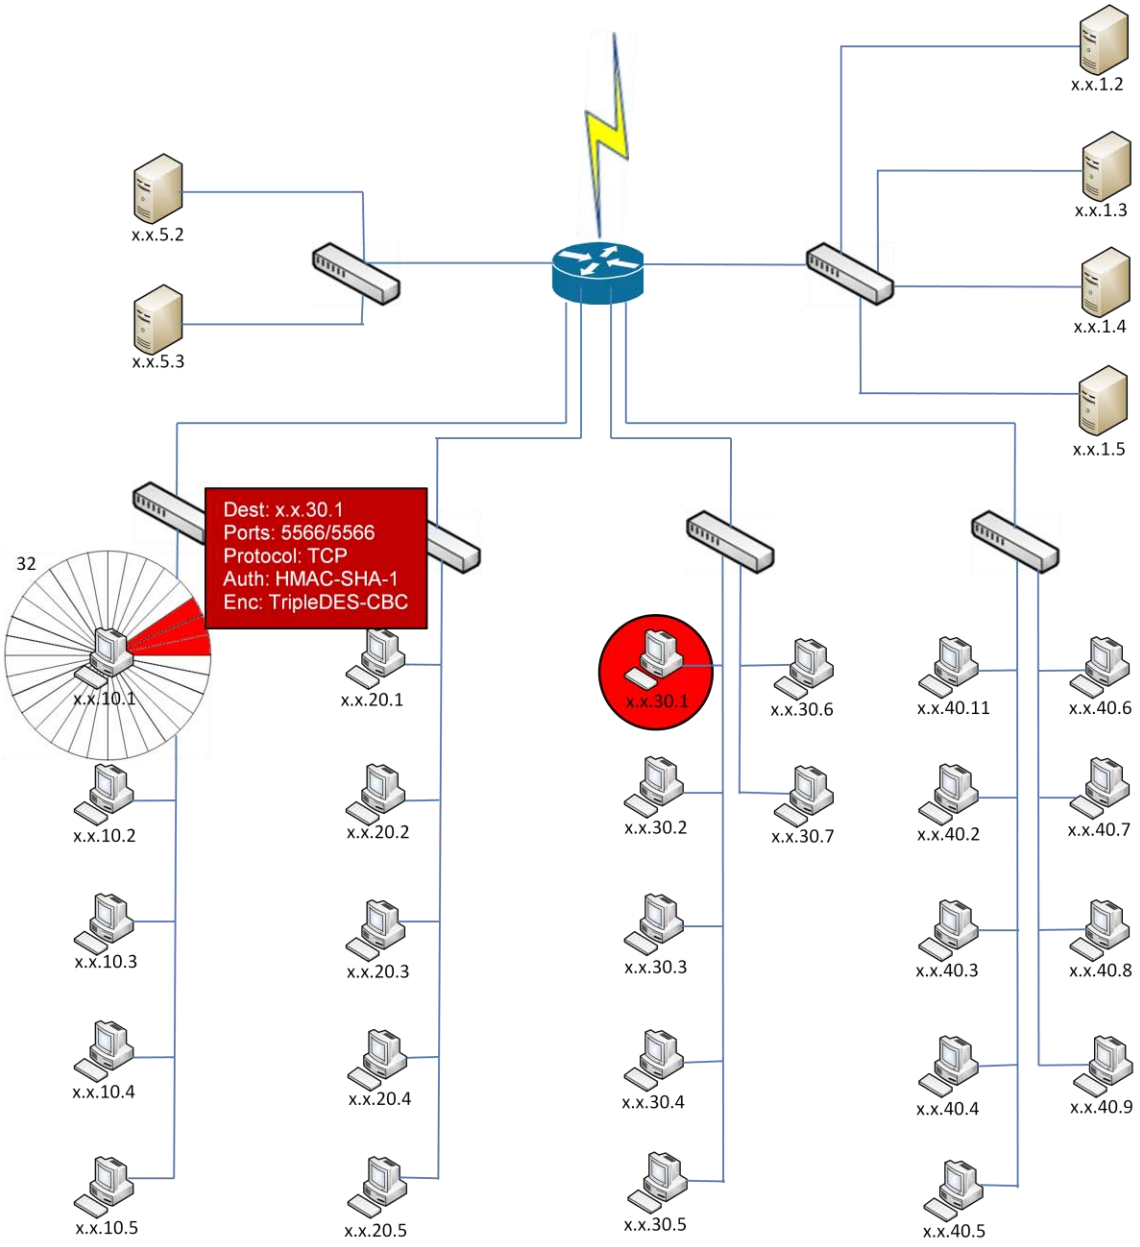

Figure 39: Using glyphs to represent IPsec data on a physical network map

4.2.3.2 *Evaluation.* This approach does not do as well representing all of the data dimensions. While the glyphs used here can represent each rule on a given node, they cannot effectively represent any additional. This means more data points associated with each rule needs to be provided in a secondary window. The data representations are straightforward but become increasingly difficult to view as the number of rules represented within a glyph increases. The physical map provides a familiar setting for administrators to work with, but the actual IPsec visualization still needs to be learned. However, due to the limitation of data points that can be displayed, there should be little cost involved with adapting to the interface. Theoretically, the glyph could be applied to an unlimited number of nodes, but the space requirements and limits to the amount of data that can be immediately presented negatively impacts the scalability of this approach. These restrictions also have a stronger impact on the ability to locate specific pieces of information quickly, though the advantage of having the physical map might be enough of a counter when searching for specific nodes. When all these points are considered, this approach does not seem well suited for managing IPsec on a production network.

It is worth reemphasizing here that there are no real restrictions on how to represent data using glyphs. Theoretically, any shape or object can be used. With further research into representing IPsec data using glyphs, an alternate schema could potentially be developed that dramatically changes the overall effectiveness of the approach.

4.2.4. Redrawing the Network Map

An important aspect to recognize in visualizing IPsec rules on a network is the usefulness of a logical map versus a physical one. Consider Figure 40.a, which depicts the physical layout of a simple network of two workstations connected to a server through a switch. To visually show an encrypted connection between the workstation on the left and the server, the connection lines could be highlighted red as shown in Figure 40.b. If communications were encrypted between the server and both workstations, the result would look like Figure 40.c. However, Figure 40.c mistakenly implies that the two workstations can pass encrypted traffic between each other through the switch, which would be a valid IPsec rule. If such an IPsec rule were added, there would be no change to the physical map in Figure 40.c. Obviously, encoding the information about IPsec rules onto the physical network map in this fashion could lead to confusion. As more nodes are added to represent a more realistic local area network, the lines of communication would increasingly overlap making it more difficult to determine what information pertained to which endpoints.

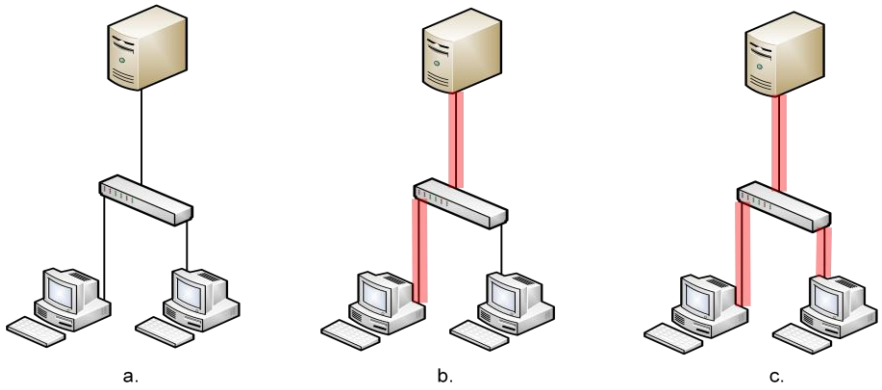

Figure 40: Example of encoding IPsec rule information on a physical network map

IPsec is concerned with end-to-end connectivity. It considers the path between two endpoints a protected communications channel regardless of the components that make up that channel. The best way to represent this is through the logical connections between endpoints. Figure 41.a shows that each workstation has an encrypted channel to the server but no connection between each other. IPsec is not concerned with the fact that the data may pass through a switch, so it has been removed from the map. In the context of IPsec, it is understood that each edge represents a single logical connection. That is to say that in Figure 41.a, there is a distinct connection represented between the server and each workstation and not a connection between the two workstations that passes through the server. When an IPsec rule between the two workstations is established, it is clearly illustrated by a connection between the two endpoints. Figure 41.b shows an example of a channel without encryption being established between the two workstations. Drawing the network map as logical connections enables the radial graph approach presented in the next section.

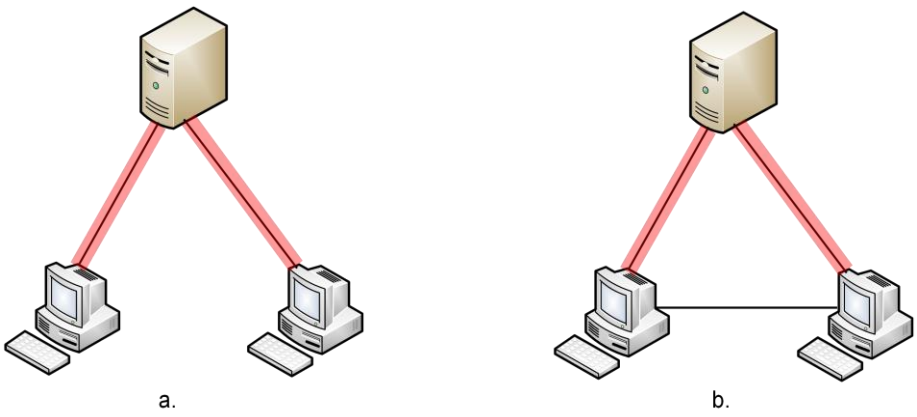

Figure 41: Example of encoding IPsec rule information using a logical network map

4.2.5. Radial Graphs

The goal to create a visualization that is intuitive to network administrators limits the number of visualization techniques that make suitable choices for visualizing IPsec rules. The decision to represent IPsec rules as logical connections between nodes further narrows the field of potential candidates and design decisions.

4.2.5.1 *Visualization*. A radial graph, similar to the Facebook Friend Wheel shown in Chapter II, is used to display the IPsec rules as logical connections between endpoints. A radial graph was chosen because it could theoretically display an infinite number of nodes while allowing a distinct line to be drawn between any two nodes or points on the circle. This satisfies the need to be able to represent both large numbers of nodes as well as the logical connections between communication endpoints. However, effectively visualizing large numbers of nodes is still an issue to be addressed. Figure 42 illustrates a view of network nodes displayed on a radial graph encoded with IPsec rule information.

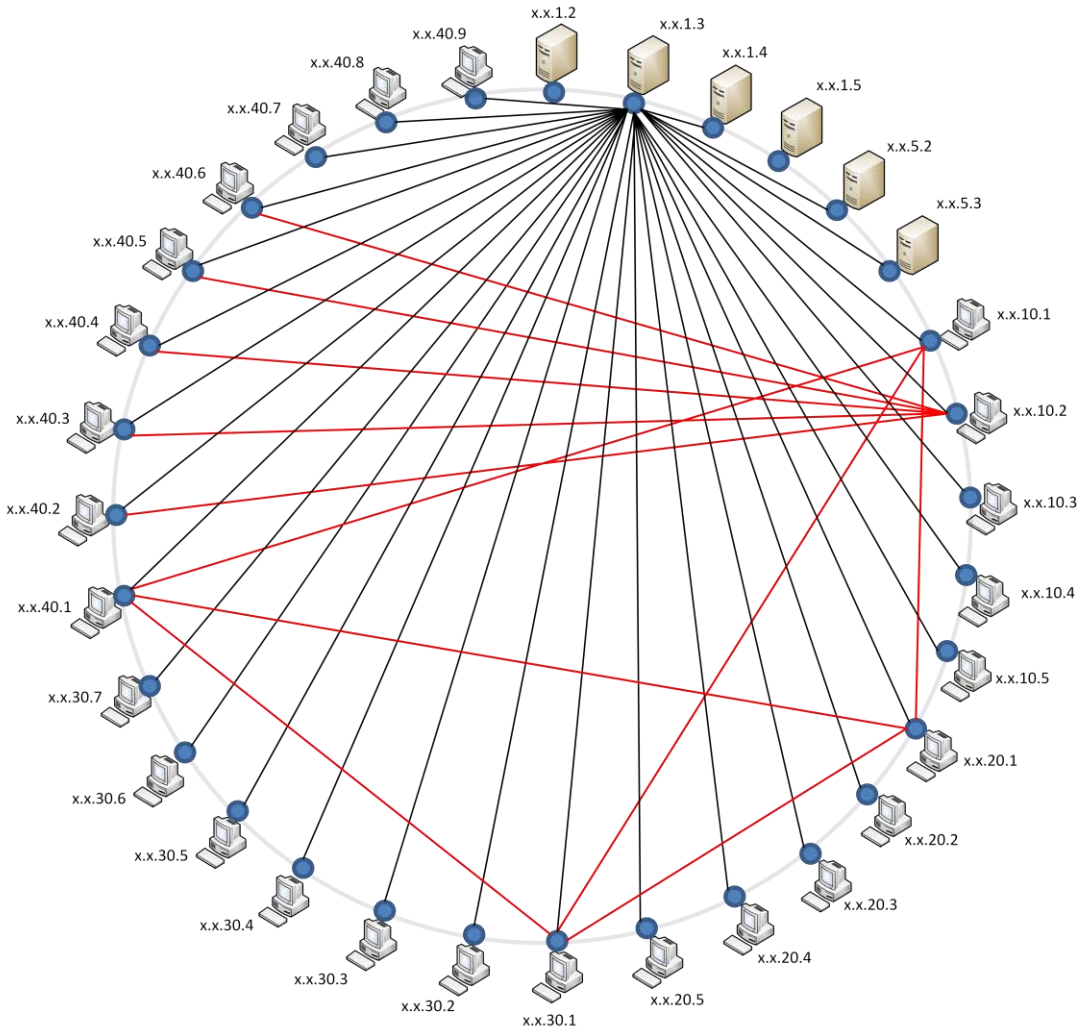

Figure 42: Radial graph showing IPsec rules as logical connections between network nodes

Regarding the representation of nodes in Figure 42, it is worth noting that the figure was produced manually rather than using the Prefuse Visualization Toolkit shown in the following sections. Experiments with various renderings led to the combination of a dot, an icon, and a label. Originally, no dot was included. Using only an icon with a label led to confusion as the icons tended to obscure some edges and endpoints. The dots worked well, but the icons were still desired for the additional information they

conveyed. The labels were essential from a management perspective, although they could be modified to include additional information such as system name. This configuration was not used in the Prefuse simulations. Upon developing an actual visual IPsec management application, a configuration for representing nodes such as the one used in Figure 42 should be considered.

4.2.5.2 *Evaluation.* The radial graph approach does well when evaluated against the criteria outlined in Section 3.2.2.1. Regarding representing all data dimensions, the display immediately identifies the source and destination address for all IPsec connections of a specific type and which rules provide encryption. Additional information that cannot be practically displayed at all times such as the algorithms associated with each rule can be displayed in a window when the cursor is placed over a link or node. It is simple to distinguish between each data representation. Using the radial graph approach, the most likely source of confusion comes from a high concentration of rules where some end points might become cluttered. This is typically unavoidable with any visualization but can be mitigated with techniques such as zooming or highlighting that are discussed in later sections. This approach is likely to be the most intuitive to network administrators of the approaches explored here. There is a clear distinction between two types of nodes (likely servers and workstations) and two types of connections represented by red and black lines. An administrator could quickly identify the connections between all 32 nodes on this LAN. Given the context of IPsec and some additional information appropriate to this example, such as the IPsec rule being observed, an administrator could now identify all network nodes implementing the specified rule

and which channels provide encryption (represented by the red lines) versus unencrypted channels. The visualization can be scaled to include a large number of nodes, but the techniques for addressing scalability such as panning and zooming become more important as the number of nodes increases. Specific information can be found quickly by selecting networks, nodes, or types of IPsec rules. As a result of these factors, the radial graph approach seems well suited to managing data on a production network.

4.2.6. Evaluation Summary

Table 7 presents a preliminary ranking of each visualization technique based on the criteria presented in Section 3.2.2.1. Each approach is ranked from 1 to 4, with 4 being the highest. The ranking is based on the preliminary evaluation presented in conjunction with each visualization approach and insights gained while developing the visualizations for the data provided.

Table 7: Evaluation summary for each visualization approach explored

	Parallel Coordinate Graph	Treemaps	Glyphs	Radial Graphs
Does visualization represent all data dimensions?	3	2	1	4
Are data representations easily distinguishable?	3	2	1	4
Is the visualization intuitive to network administrators?	2	1	3	4
Is the visualization scalable?	4	2	1	3
Can specific information be found quickly?	3	2	1	4
Is the approach suitable for managing data on a production network?	3	1	2	4
Total score:	**18**	**10**	**9**	**23**

Clearly the radial graph approach appears to be the best candidate for further evaluation as a management approach to IPsec. As a result, we developed a simulation of a potential interface for an IPsec management tool using the Prefuse Visualization Toolkit. The interface and potential functionality are described and presented in detail in Section 4.3.

It should be restated that the evaluation presented in this section represents only a preliminary evaluation based on personal network management experience, visualization research, and insights into each approach gained while developing the visualizations. To further develop a tool like the one presented in this thesis, these results should be validated through additional development of the visualizations and a usability study. If possible, inputs should be gathered from visualization and network management professionals as well as individuals with varying degrees of computer experience.

4.3. Using Visualization for IPsec Management

Developing visualizations for IPsec rules is only part of the puzzle. The visualization must enable efficient management of IPsec on a network. This section suggests a common layout for integrating one of the visualization schemas developed in the previous section into a management tool. It presents various potential views and management features to explore the viability of the visualization as an approach to IPsec management.

4.3.1. Interface.

An interface that would be appropriate for a network management tool and familiar or intuitive to a network administrator is desired. Therefore, a layout common to

computer applications and used in some existing network management tools is adopted here. Figure 43 shows a screenshot of LANsurveyor [35], a network mapping tool available from SolarWinds. The interface provides toolbars across the top presenting icons for common actions, a panel down the left side for selecting operations, displaying information, etc., and a main window for displaying the network map. The screenshot shows the ability to right click a node to select specific operations as well. There is also a small window behind the right-click menu showing an overview of the map with a movable window for navigation.

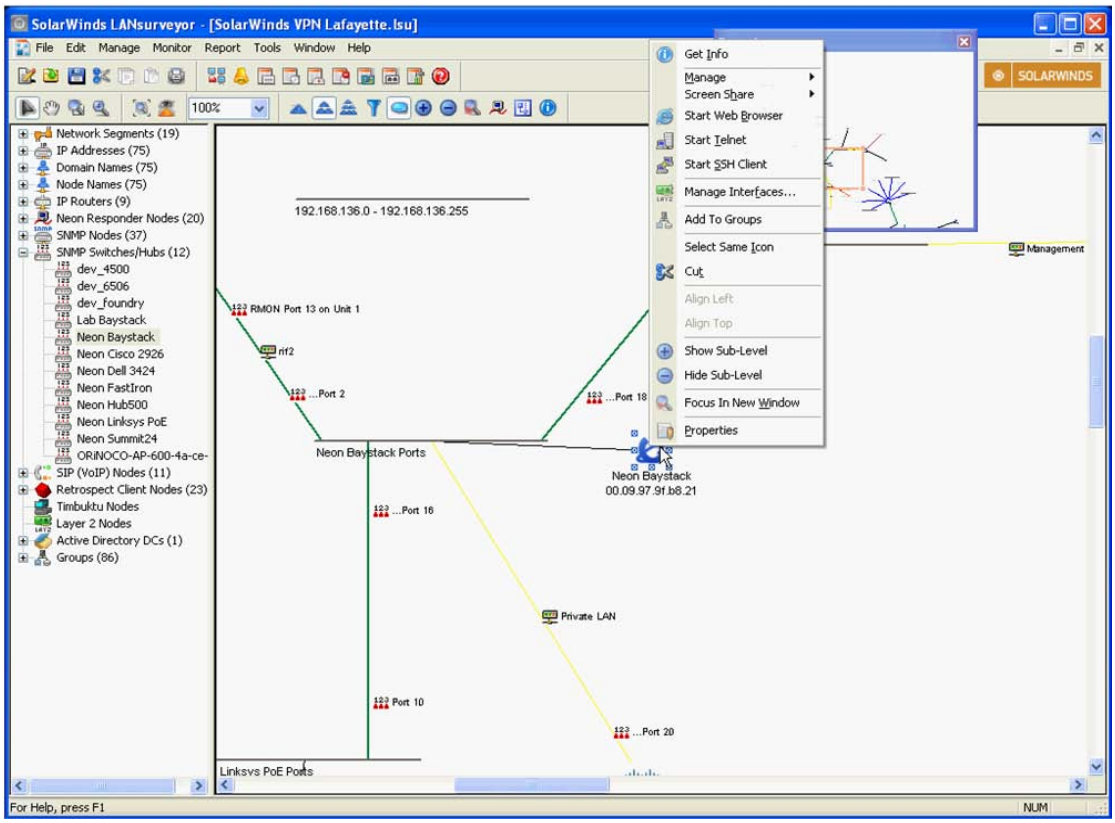

Figure 43: LANsurveyor network diagramming tool [35]

Combining the common layout features illustrated in Figure 43 with the use of radial graphs to depict the logical connection between nodes produces an interface similar to the one shown in Figure 44 below. There is a toolbar across the top for implementing menus or quick launch buttons. The left panel provides space to display the rules that have been deployed on the network and selectors for additional functionality or information to be displayed. The main window provides space for the visualization to be displayed, and interactive features can be incorporated such as highlighting endpoints and having an information pop-up window appear when placing the cursor over a link as seen in Figure 44.

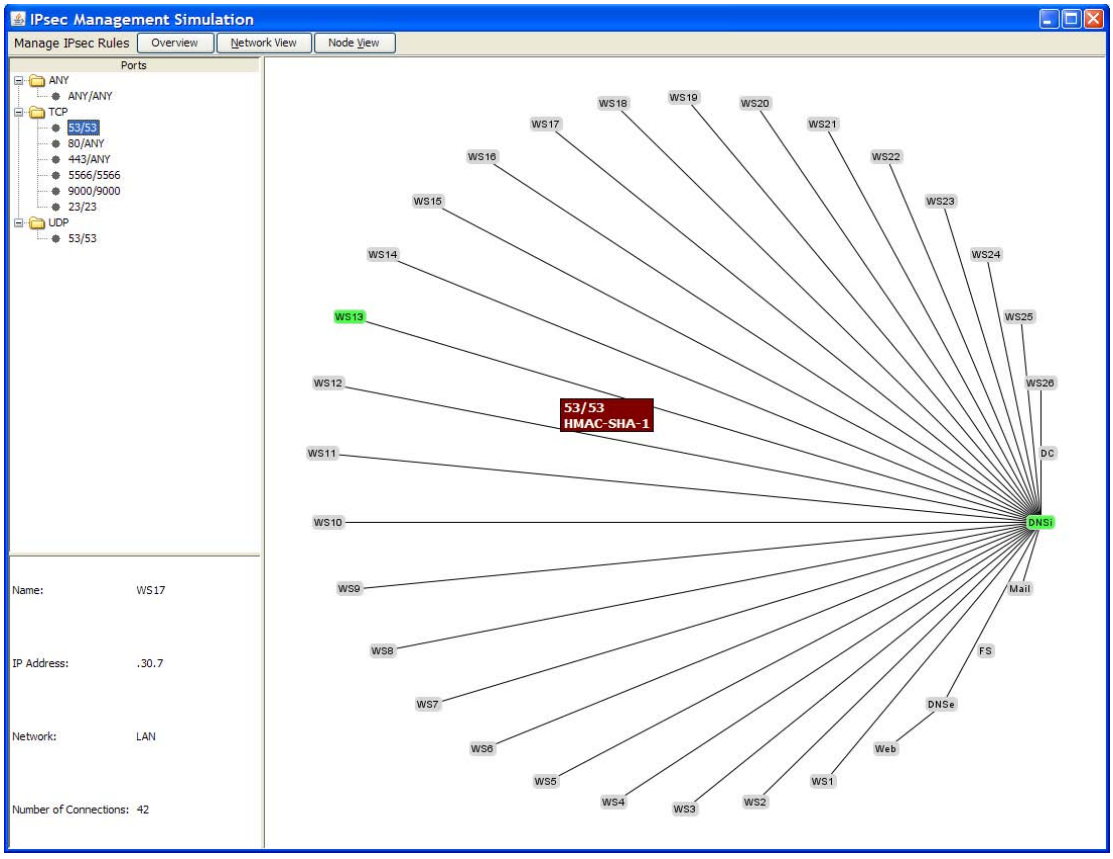

Figure 44: Sample interface for visual IPsec management tool using radial graphs

4.3.2. Views

Using the interface suggested in Section 4.3.1, this section explores how the visualization might be used by an administrator to actually manage IPsec rules deployed on a local area network and between a LAN and external networks. It is important to note that many other interfaces are possible. The one presented here is used to explore the viability of using visualization to manage IPsec.

4.3.2.1 *Overview*. When an administrator chooses to manage IPsec, there needs to be a logical place to start. There may IPsec connections to nodes both internal and external to the administrator's network. Since policies are likely different on each external network an administrator's network shares connections with, it makes sense to separate administration of rules found solely on the administrator's network and the rules shared between that network and each external one.

Figure 45 suggests a potential starting point that would allow an administrator to select which area to manage. An icon is used to represent the local network and any external networks the local network shares any IPsec connections with. No specific IPsec information is presented on this screen other than an indication that IPsec rules exist between the local network and any external networks displayed. There is, of course, the potential to encode additional information if desired. This could include information such as the number of nodes on each external network with IPsec connections to the local network or the total number of IPsec connections between the local and external networks. Additional information was not included for the example here since much

more detailed information about the rules becomes more evident when viewing specific areas.

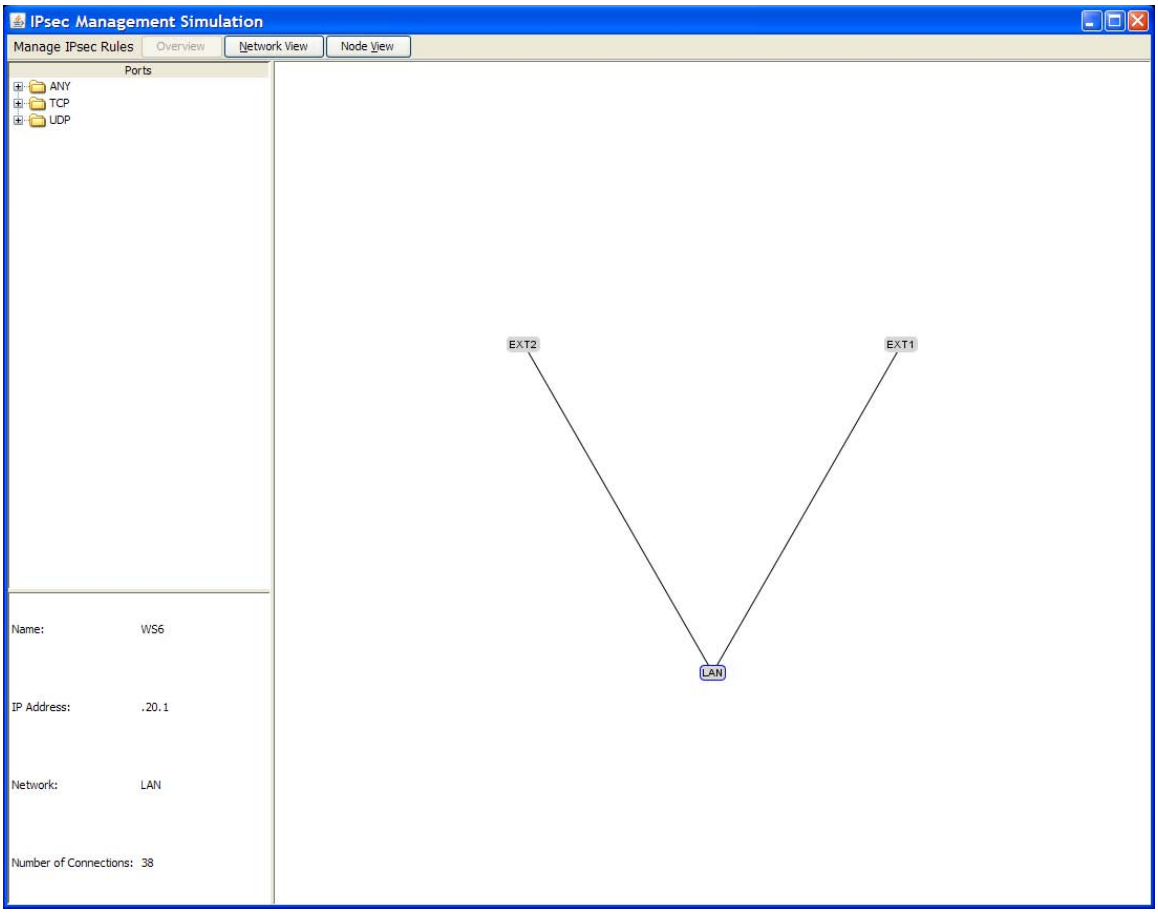

Figure 45: Overview screen to allow an administrator to choose between managing local or external connections

Double clicking on an external view could display the connections shared between the selected external network and the local network. Double clicking on the local network could likewise bring up a detailed display of rules shared between local network nodes. We first examine potential views within the local network.

4.3.2.2 *Network View*. Selecting the local network displays all local network nodes in a radial graph. An administrator identifies which rules to display by selecting a type from the left panel. Figure 46 shows all IPsec rules for HTTP traffic within the local network.

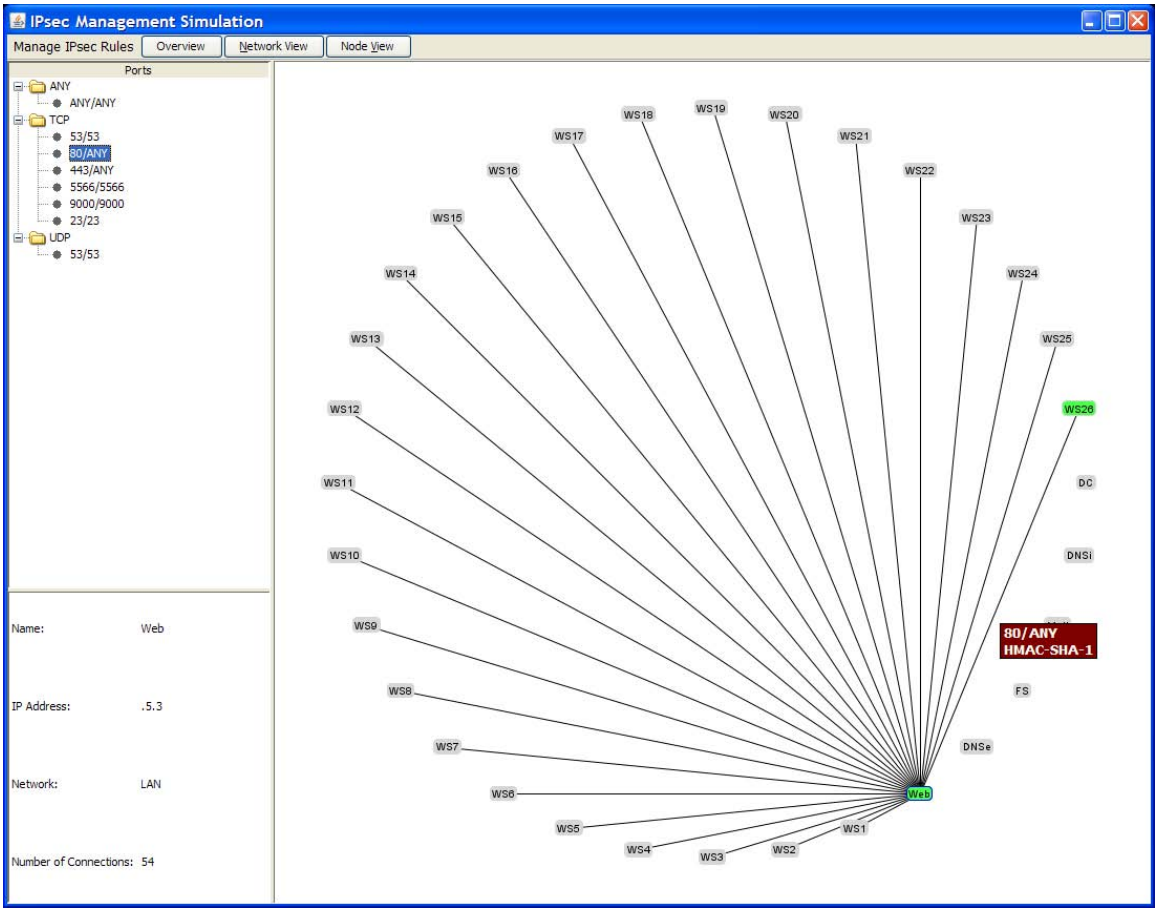

Figure 46: View of all network nodes with IPsec rules for HTTP traffic

The network nodes without HTTP rules still show on the radial graph. They serve as a visual cue to an administrator as to where a particular rule is not being used, which may be just as important to know as where the rule is being used. Selecting another rule type from the left panel has no affect on the nodes, but will usually change the edges of

the graph. Only in situations where different rule types provide the same protections between the same sets of nodes would the display remain completely unchanged. Figure 47 shows how the display would look if an administrator selected HTTPS from the list of rule types.

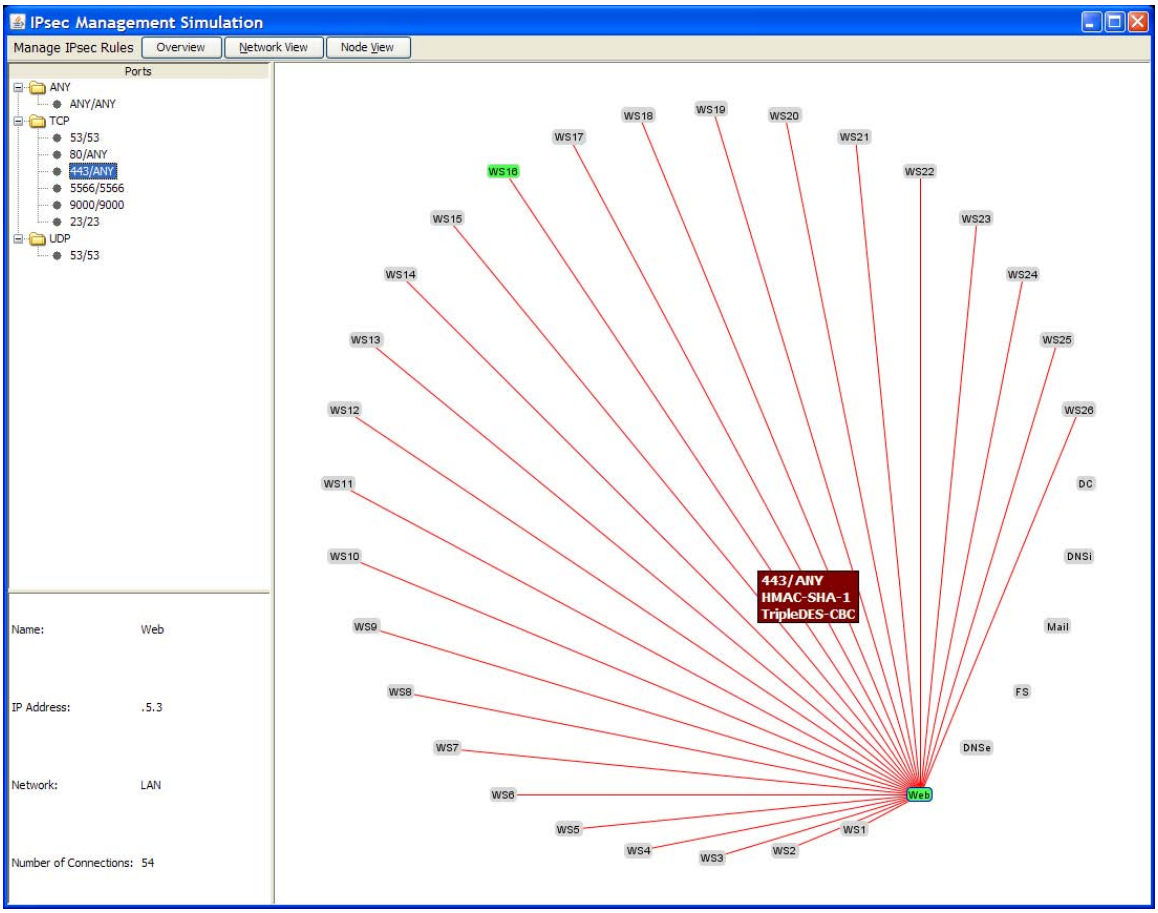

Figure 47: View of all network nodes with IPsec rules for HTTPS traffic

As you can see, the two views are the same except the edges in the display for HTTPS traffic in Figure 47 are red to indicate that IPsec provides additional encryption to HTTPS traffic. The edges are otherwise unchanged because there are rules for both HTTP and HTTPS traffic between the web server and each workstation on the network.

In addition to being able to view the IPsec rules for all nodes on the network simultaneously by type, an administrator might want to manage the IPsec rules applied to a specific node. The following section how managing a single node might look using the suggested interface.

4.3.2.3 *Node View*. Managing the IPsec rules on a single node is mostly the same as with the view of the entire local network. However, there are two important differences. First, there may not only be nodes on the network that do not have a specified rule type in common with the selected node, but there may also be nodes that have no IPsec connections shared with the selected node at all. Second, even though we are selecting nodes internal to the local network, the individual nodes may have IPsec connections with nodes found in external networks. To address the first issue, consider Figure 48, which shows the display as it looks after double clicking on the web server node in Figure 47.

Through a simple animation, the web server node moves to the center of the graph and the edges from Figure 47 are redrawn. Notice the external DNS node (DNSe) is located on the main ring with no connection drawn to it. This is because the external DNS server shares an IPsec rule with the web server but it is not an HTTPS rule, which is the type specified in the display. However, the domain controller (DC), internal DNS (DNSi), mail server (Mail), and file server (FS) nodes have been moved to an outer ring. This indicates that these nodes have no IPsec rule between them and the web server. They are included in the display as visual cues to the administrator. A node that requires the selected rule that is missing it, for example a new workstation that has been

improperly configured, would show up on this outer ring alerting an administrator to take action.

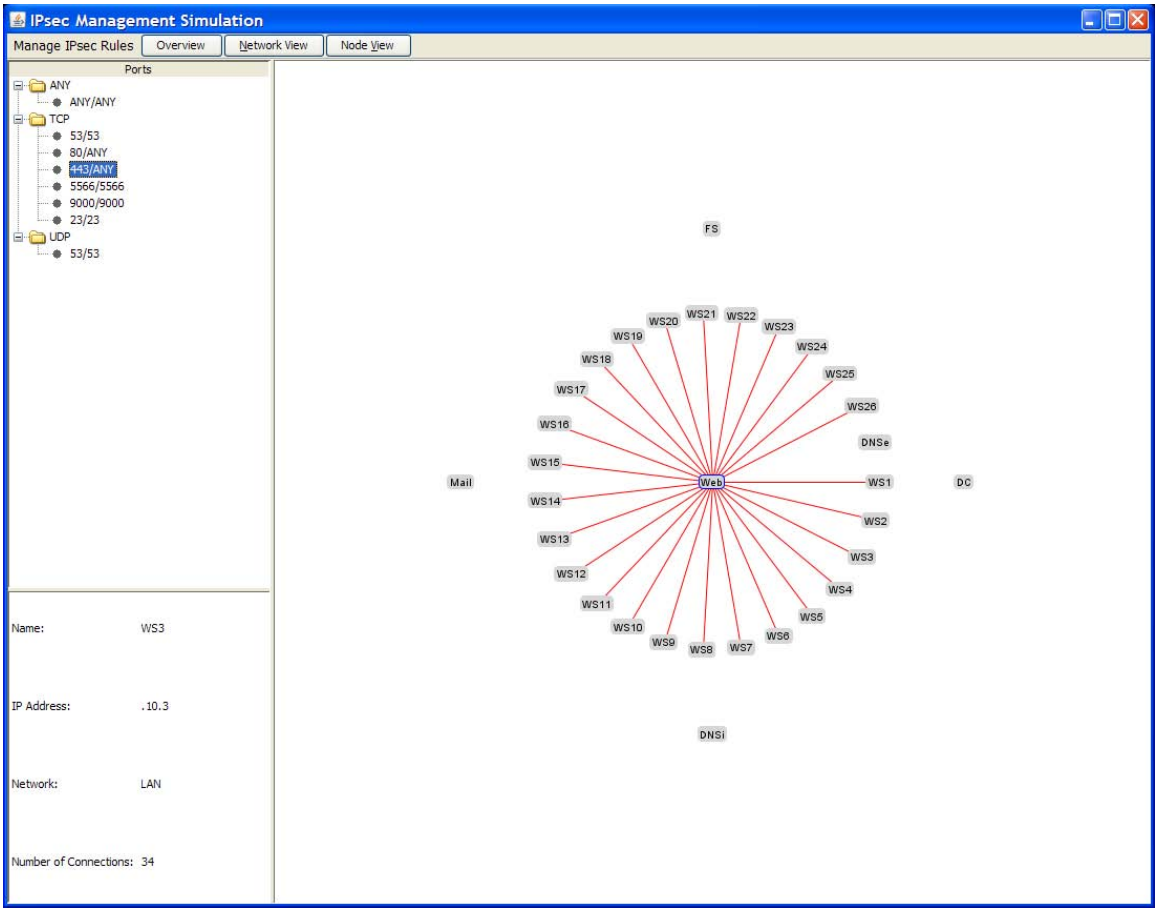

Figure 48: View of IPsec rules for HTTPS traffic found on the web server

We now consider the case where a specific node shares IPsec rules with nodes outside the local network. Figure 49 shows the changes to the display in Figure 48 after selecting the IPsec rule type ANY and double clicking on the WS1 icon for Workstation 1 (x.x.10.1).

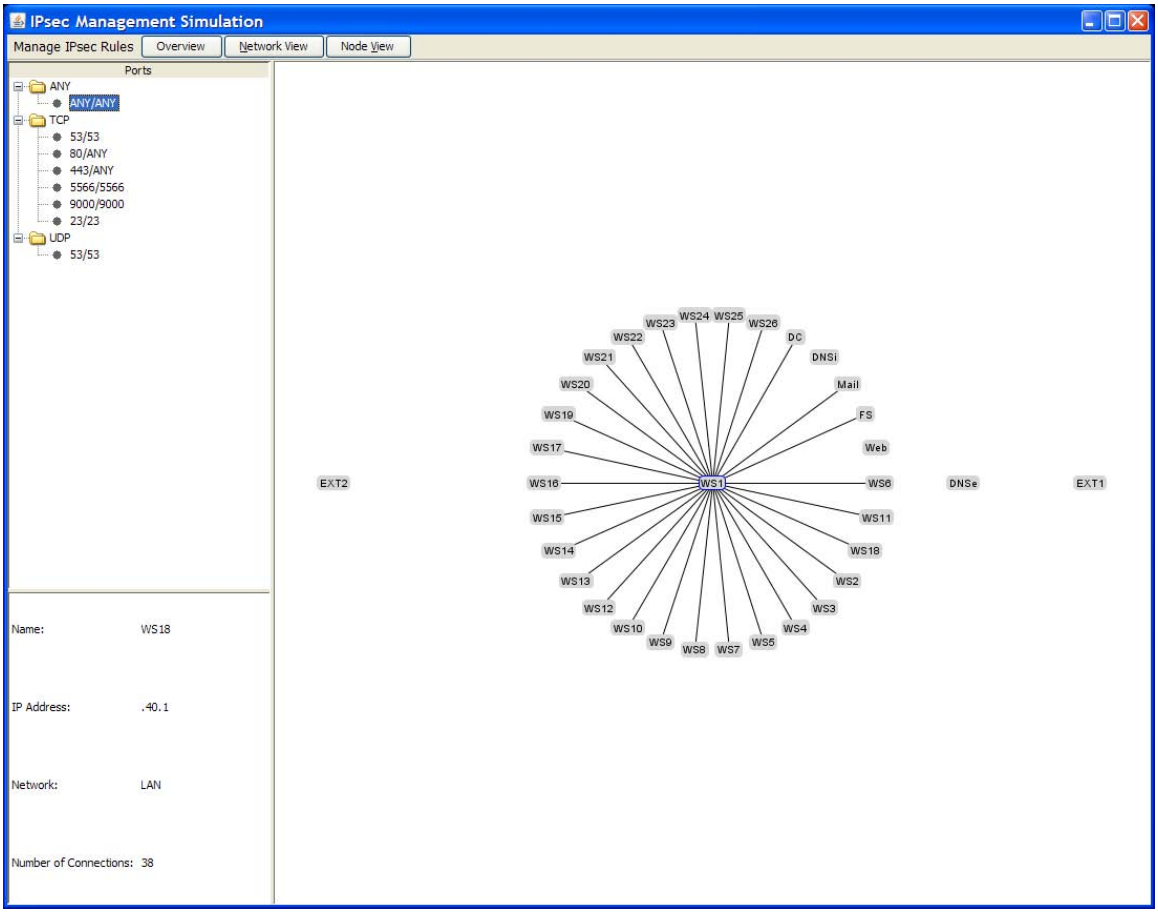

Figure 49: View of IPsec connections to a specific node including connections to external networks

Figure 49 shows the display redrawn with WS1 in the center. There is a node (Web) on the inner ring with no connection to the selected node (WS1). As before, this indicates that the web server has a connection to the selected node but it is a different type of rule than the one selected in the left pane. Also, as with the previous example in Figure 48, there is a second ring that shows nodes on the local network that have no IPsec rules in common with the selected node. However, in Figure 49 there is a third ring with two icons on it: Ext1 and Ext2. These icons identify two external networks with nodes that share IPsec connections with the selected node in the center. These single icons can

represent any number of connections to the selected node. To see the actual connection between the selected node and the nodes within a given external network we can double click on one of the external network icons.

Figure 50 shows the results of double clicking on the Ext1 icon in Figure 49. All IPsec connections between Workstation 1 and any nodes in the selected external network are displayed. In this case, nodes on the external network that do not share IPsec rules with the selected node are not displayed on any ring since the system is not aware of those nodes. Notice that the nodes for the local network have been replaced by an icon labeled LAN and placed on the outer ring.

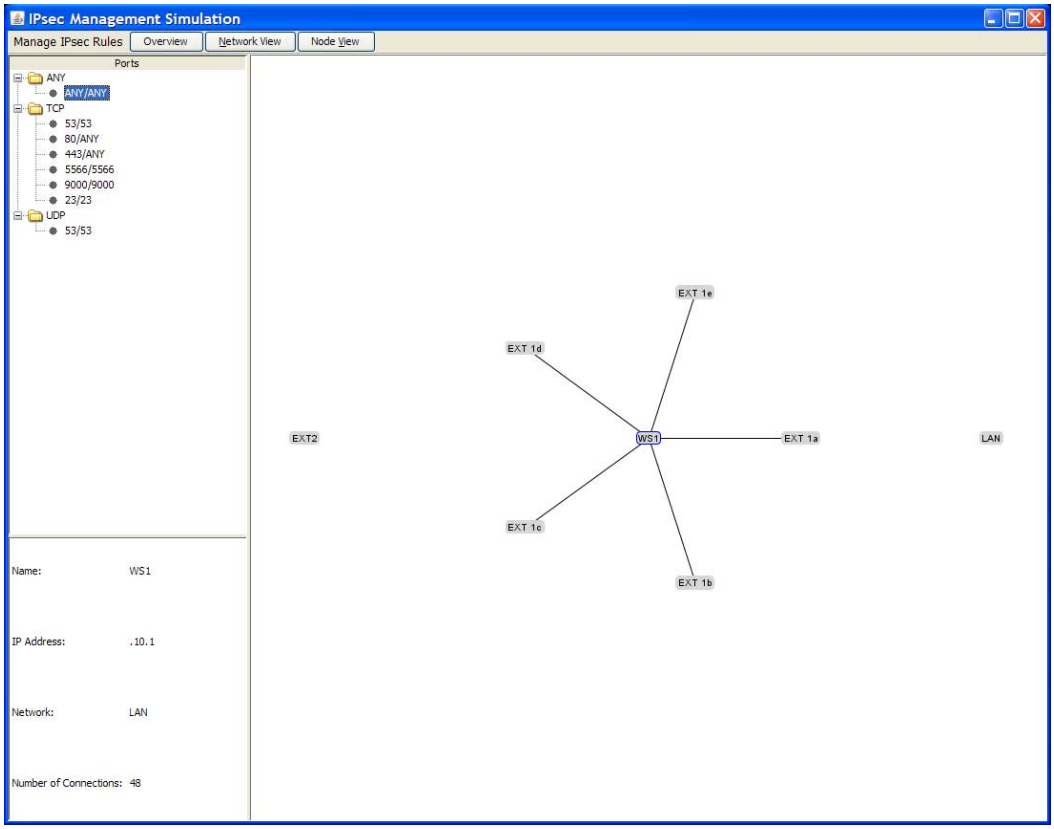

Figure 50: View of IPsec rules between a local network node and nodes on an external network

This approach allows the administrator to be cognizant of all connections at all times yet maintain focus on logically separate areas at any given time. As mentioned in Section 4.3.2.1, an administrator might initially choose to manage connections between the local network and an external network. This scenario is explored in the next section.

4.3.2.4 *External View*. The proposed overview screen in Figure 45 identifies any external networks that share IPsec rules with nodes in the local network. Double clicking on one of the external network icons would display the IPsec rules between the nodes on each network. As with other views, all nodes on the local network would be displayed for administrative purposes, even if they had no IPsec rules applied. Just as external nodes without IPsec rules shared with local nodes were not displayed in Node View, they would not be displayed here. Figure 51 provides a potential display of the IPsec connections between the local network and an external network. The nodes of each network are displayed on an arc or semicircle to maintain the overall design approach of using radial graphs.

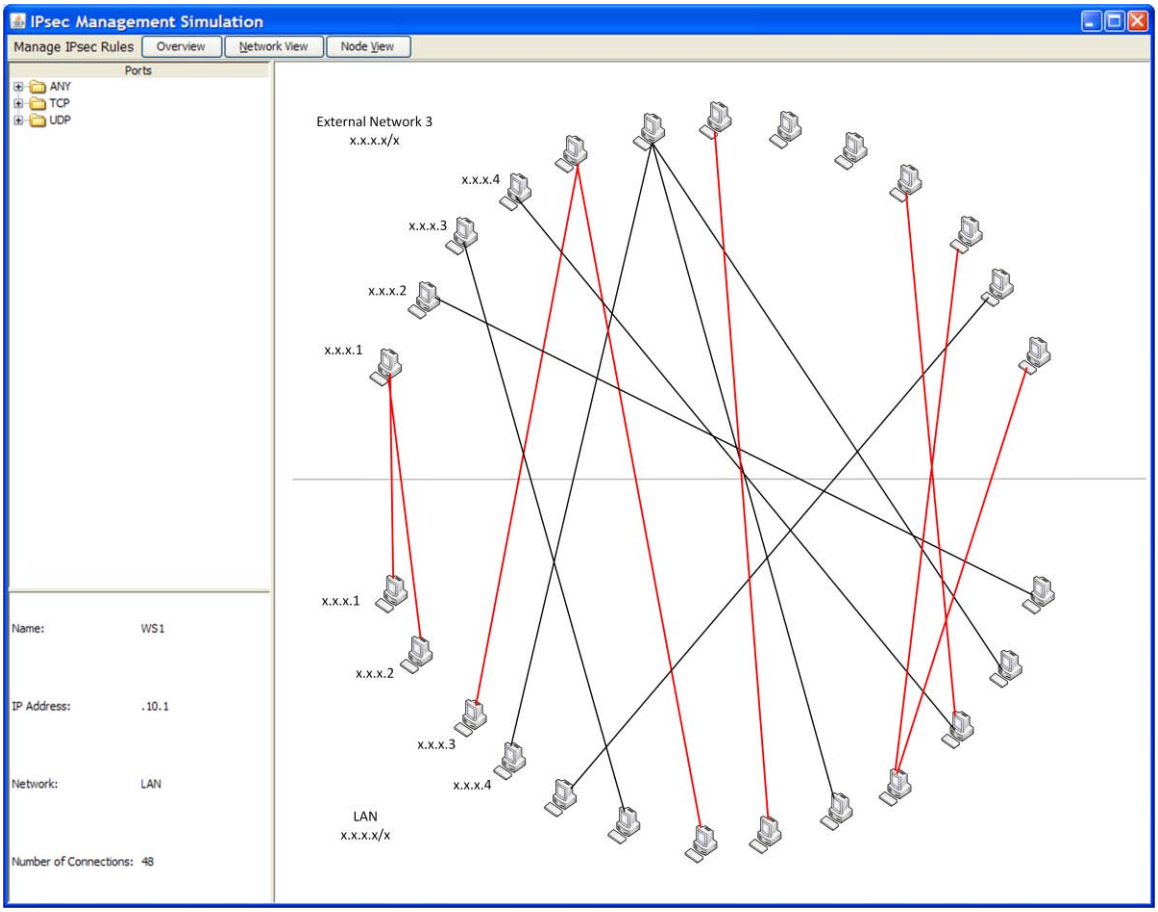

Figure 51: View of IPsec rules between the local network and a selected external network

4.3.3. Management Functionality

The goal behind finding a suitable method of visualizing IPsec rules on a network is to enable effective management of those rules through the visualization. To this point, the suggested interface has illustrated how the data could be presented to enhance an administrator's situation awareness regarding how IPsec is deployed on the network. However, management functionality needs to be incorporated as well. The screenshots presented so far have hinted at some features such as highlighting the endpoints and popping up an information box when moving the cursor over a link and the information

window in the left panel that can provide details about a selected node or link. There are many other features that would need to be available for basic management such as importing and exporting rules, importing new algorithms, adding new rules, and editing or deleting existing rules. These features would typically be accessible through a drop-down menu like the 'Manage IPsec Rules' tab seen on the toolbar in the previous screenshots or by right-clicking on a specific node. It is not practical to try to present an exhaustive list of potential features here, but we cover how the basic management functionality could be worked into the interface to provide a more complete picture of how the tool would actually be used.

To get a feel for how management functionality would fit into the proposed interface, we consider the functions to add, edit, and delete rules on a specific node. Additional functionality like importing and exporting rules and algorithms would all have similar interfaces. This type of functionality should perhaps be limited to Node View since this is where an administrator is focused primarily on managing the rules on a specific node. For our purposes, we will consider Node View with WS1 as the selected node in the center. Figure 52 shows the result of right-clicking on WS1 and selecting 'Add Rule'. A dialog box appears allowing the administrator to specify all of the information required to establish a rule. Following the scenario for a streamlined IPsec implementation outlined in this thesis, default values are included for the authentication and encryption algorithms. The default for encryption is 'NULL'. By changing the value for the encryption algorithm, the administrator is effectively indicating that encryption is required. Figure 52 shows the drop-down box for the encryption algorithms being activated.

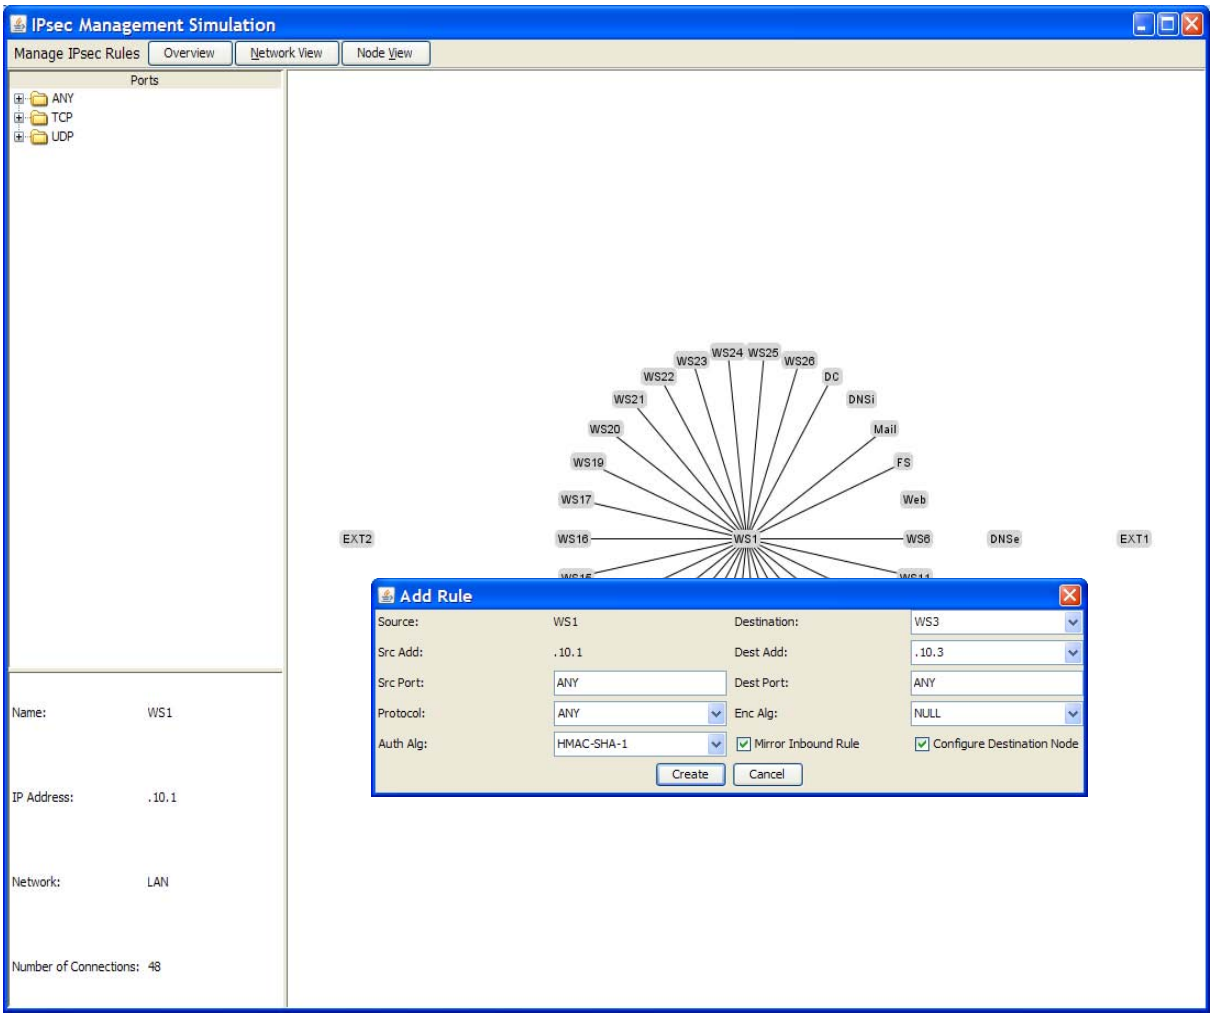

Figure 52: Add Rule dialog box

The source node information is automatically populated. The administrator can then specify the destination node, ports and protocol, and the algorithms required if they are different than the defaults. This configures the rule for traffic outbound from the source. The checkbox labeled 'Mirror Inbound Rule', checked by default, configures the inbound rule on the source node as well. There is an additional checkbox labeled 'Configure Destination Node'. This allows the administrator to have the corresponding rule created automatically on the destination node provided the administrator had the

necessary permissions to the node. This feature would be grayed out if the destination node was external to the local network.

Editing an existing rule works in much the same way as adding a rule. Figure 53 shows the result of right-clicking on WS1 and selecting 'Edit Rule'.

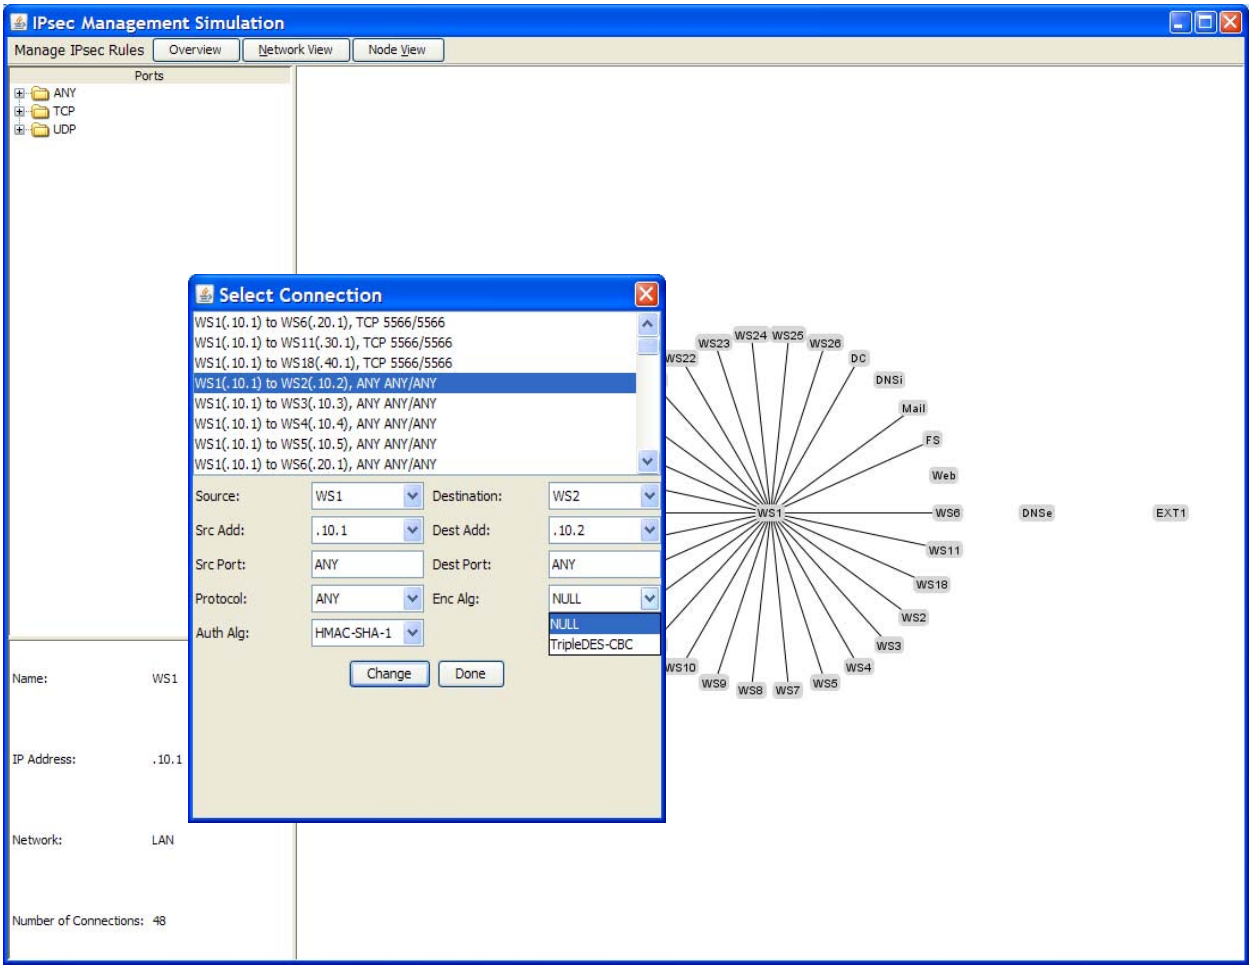

Figure 53: Edit Rule dialog box

The 'Edit Rule' dialog box presents a similar interface as the 'Add Rule' dialog box, but it includes a list of all rules configured on the selected node. Selecting one of

the rules fills in the information in each box below. The administrator can then change any of the fields as required. This feature is currently configured to make the requested changes on the inbound and outbound rules on both the source and destination node. This means that if an administrator wanted to change the configuration on any combination of these rules other than all of them, the rules would need to be deleted and new rules would need to be created. The 'Edit Rule' dialog box could easily be modified to allow an administrator to select a specific rule on either node.

The simplest of all, the 'Delete Rule' feature simply presents a list of all rules configured on the selected node. The administrator can then select a rule from the list and click 'Delete'. As with 'Edit Rule', this will delete the inbound and outbound rules from the source and destination although options could be added to allow for more granularity. Figure 54 shows the result of right-clicking on WS1 and selecting 'Delete Rule'. For all three functions, the administrator is prompted to confirm the action before continuing.

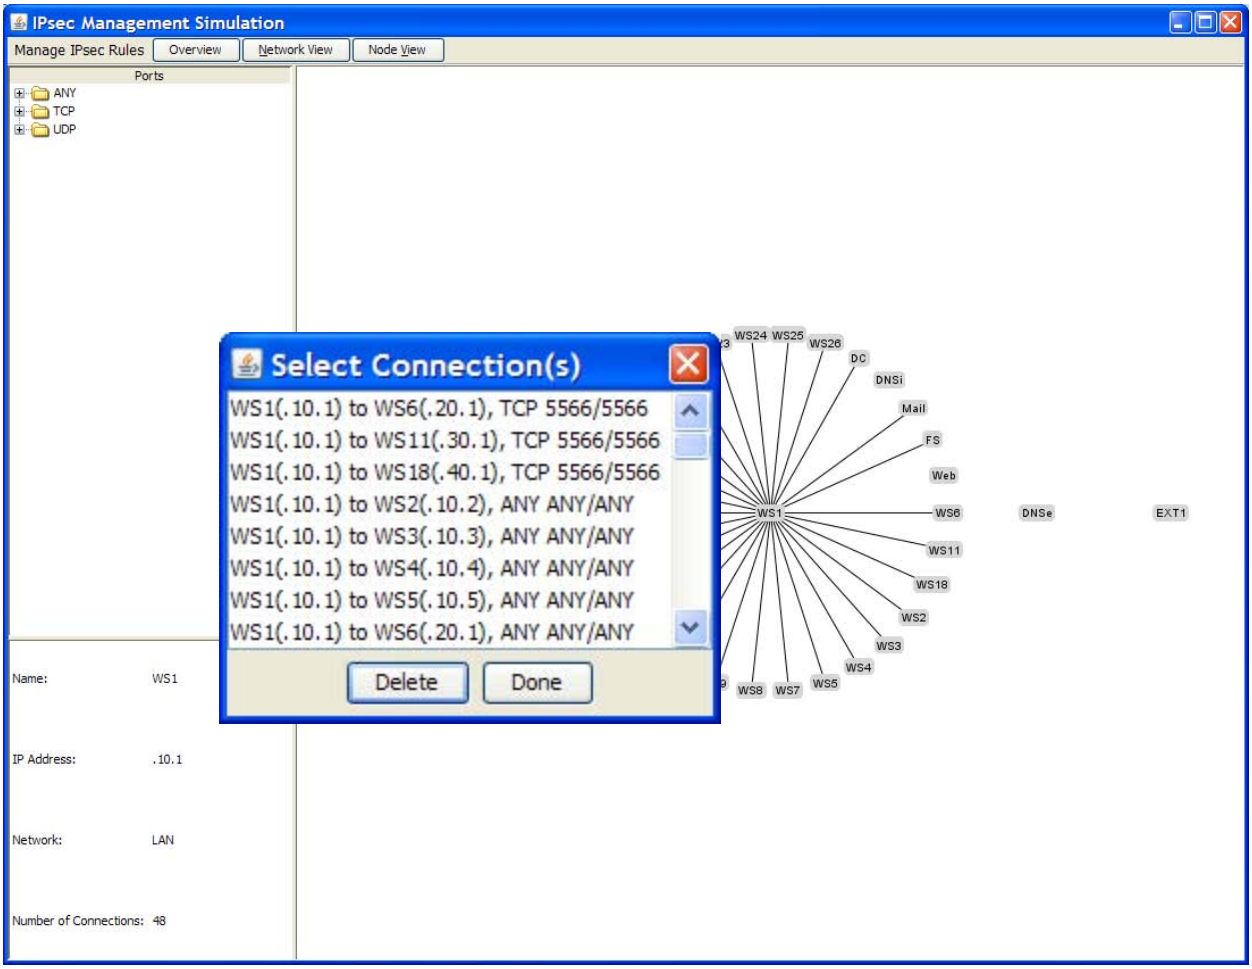

Figure 54: Delete Rule dialog box

4.4. Summary

The results presented in this chapter firmly pave the way for a new direction in IPsec management. We explicitly qualified the potential impact simplifying IPsec could have on managing it, which has seemingly only been hinted at throughout IPsec literature. We presented original visualizations of IPsec rules using various techniques and also illustrated the effects streamlining IPsec would have on visual representations of the data. Also, some novel approaches to existing visualization techniques were explored to adapt them to effectively represent the data. Each technique was evaluated to

determine its viability as an approach to IPsec management, and one approach was selected for further exploration in a simulated IPsec management application. This simulation provided hands on experience and valuable insight into how the selected visualization approach could actually be used to manage IPsec rules on a production network. Throughout the various evaluations, many issues surrounding functionality and scalability were explored. The result is a solid foundation for developing a robust IPsec management tool that could be adapted to fit other similar systems as well.

V. Conclusions

5.1. Research Goals

The goal of this research was to explore the viability of simplifying IPsec management through visualization in a way that was intuitive to network administrators and applicable to any system of logical connections between networked nodes. To do this, we first explored steps to simplify the IPsec implementation and their effects on IPsec administration and the visualization of IPsec rules. Additionally, we applied several visualization techniques to a dataset of IPsec rules to evaluate different approaches and gauge their effectiveness in representing IPsec rules. Ultimately, this allowed us to incorporate a visual representation of IPsec rules into a simulated IPsec management application to explore its potential for managing IPsec on a production network.

5.2. Results

Chapter IV details the results of each phase of this research effort. The sections below provide a summary of those results for each area in Chapter IV that show how our research goals were met.

5.2.1. Simplifying IPsec

Regarding the IPsec implementation, we were able to explicitly show the potential effects of simplifying administration by reducing the number of options available to an administrator when configuring IPsec rules. The overall result is a work load reduction of approximately 45% when configuring the common case rule. More importantly, we

were able to show the effects that simplifying the implementation would have on visual representations of IPsec rules.

The decision to eliminate one operating mode and one protocol eliminated the need to represent not only those two data points, but combinations of those data points as well where, for example, the Authentication Header (AH) protocol is used to provide authentication while the Encapsulating Security Payload (ESP) protocol is used to provide encryption over the same channel. Additionally, the decision to always provide authentication eliminated the need to distinguish between three potential configurations where authentication only, encryption only, or both were provided. Figure 55 shows the impact of these decisions on a visualization of IPsec rules using the radial graph approach.

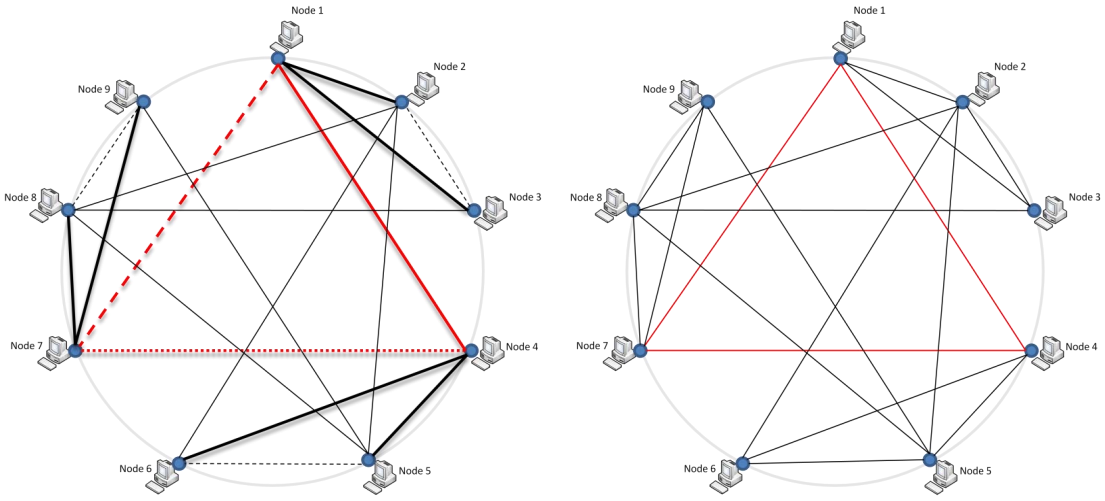

Figure 55: Side-by-side comparison of radial graph visualization of sample IPsec rules without streamlining decisions (left) and with streamlining (right)

These decisions resulted in cleaner, simpler visual representations of the IPsec data. Regardless of the visualization approach, more information could be provided in the same space with less clutter providing a more efficient presentation of the data.

5.2.2. Visualizing IPsec Rules

With regards to visually representing IPsec rules, we were able to develop and evaluate visualizations using four widely varying techniques. These visualizations provided insights into how IPsec rules could be represented and allowed us to evaluate how effectively and intuitively the data could be presented to an administrator.

Two approaches, parallel coordinate graphs and treemaps, allowed for a large amount of data to presented in a confined space. Both were highly scalable, though as with most data visualizations, the visualizations become more cluttered as more data is represented. This, however, prompted experimentation with some novel applications of these approaches to increase their viability. What both approaches lacked was the look and familiarity of existing network management tools that would help make them intuitive to administrators.

To address familiarity, glyphs were applied to a physical network map that might be used for network management to encode IPsec information onto a visualization that network administrators might be immediately more comfortable with. While this approach had potential, further research to find the most effective glyphs and their application to a physical layout is necessary to determine the true applicability of the approach.

Finally, we rendered the IPsec rules using radial link-node graphs as shown in Figure 56. This allowed individual nodes to be represented using icons and connections to be shown using lines, or edges. Even though each edge represented a logical connection between two nodes, this layout presented the data in a more intuitive way to network administrators than the other approaches.

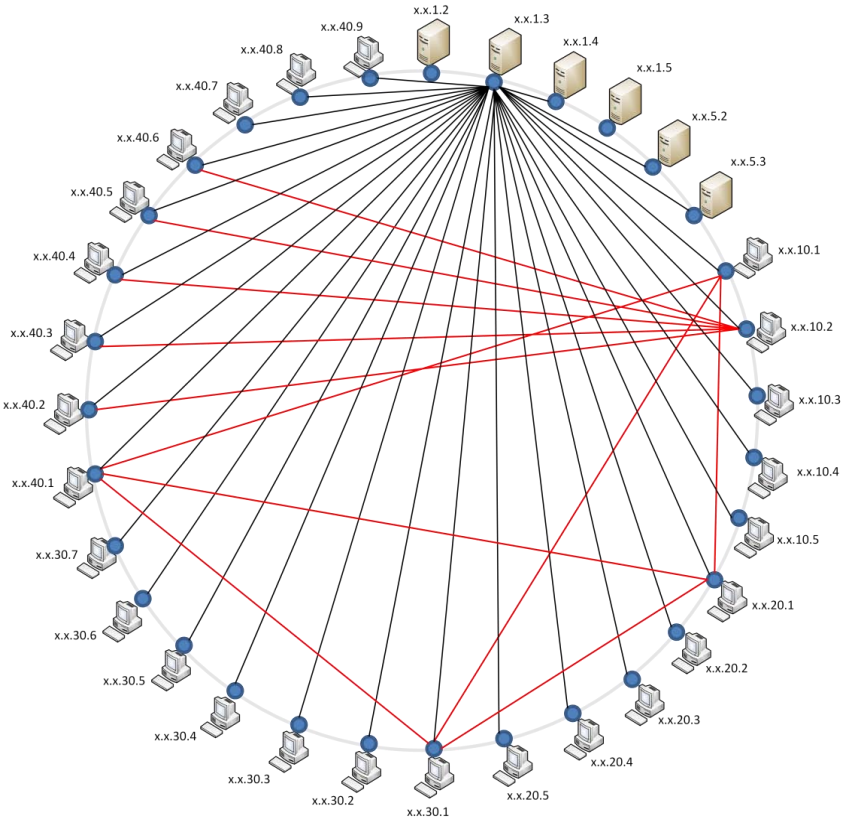

Figure 56: Radial graph showing IPsec rules as logical connections between network nodes

It was not practical to render the data for all rules simultaneously using this approach because overlap caused confusion and too many rules rendered the visualization unreadable. However, the parallel coordinate graph and treemap approaches showed that

even if the data could be presented simultaneously, from an administration perspective it was only practical to manage a single node or rule at a time. Since each radial graph was a representation of either all nodes sharing a specific rule or all of the rules on a single node, the radial graph approach seemed to be the best choice in exploring the viability of using visualization as the foundation of an IPsec management tool.

5.2.3. Using Visualization for IPsec Management

Simply providing a visual representation of IPsec rules would not have been sufficient. To evaluate the viability of using visualization to manage IPsec, we incorporated the radial graph visualization into a simulated IPsec management tool developed using the Prefuse Visualization Toolkit as seen in Figure 57. We provided IPsec rules using a simple table as they might be stored in a Security Policy Database. This enabled us to quickly modify the dataset when needed by altering the entries in the table. The simulation also allowed us to experiment with the interface, evaluate potential functionality, and identify potential issues dynamically far beyond what would have been possible using static images. As a result, we have laid a solid foundation upon which a powerful IPsec management tool could be built.

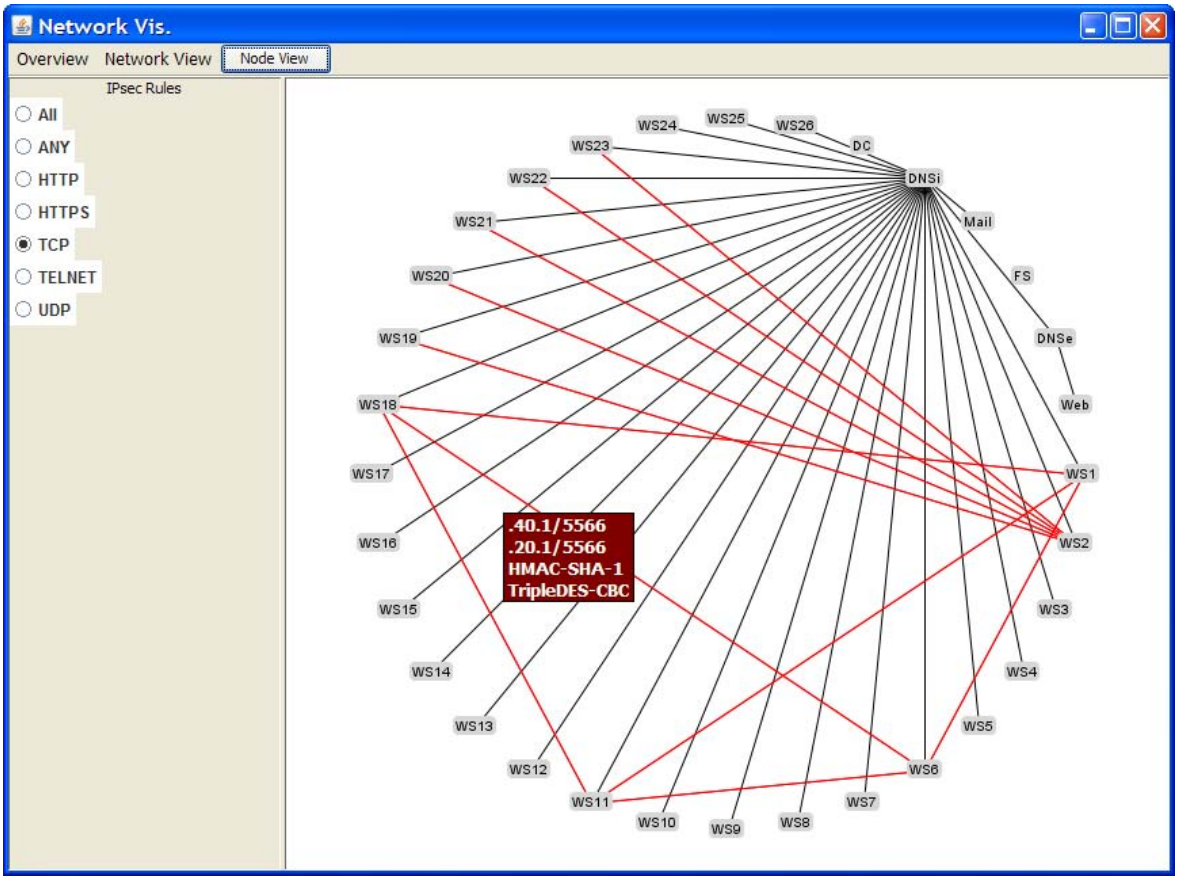

Figure 57: Screenshot of interface using radial graph visualization of IPsec rules developed using the Prefuse Visualization Toolkit

5.3. Research Contributions

The contributions of this research stem from the application and evaluation of various visualization techniques in relation to simplifying the management of IPsec or any similar system of logical connections between networked nodes. This research resulted in several original visual representations of IPsec data and novel applications of some established visualization techniques. It also explicitly illustrated for the IPsec community the effects that simplifying some aspects of the IPsec protocol suite could have on administration, including the visualization of IPsec data. Finally, this research

laid the groundwork for building an IPsec management application that can leverage the benefits of visualization to ease some of the burden of IPsec management and improve an administrator's overall situation awareness.

5.4. Future Work

As stated in Chapter III, producing a full-featured IPsec management application like the one suggested in this thesis would involve all of the many aspects of system development. Though this additional work is significant, the suggestions here are aimed toward developing the visualizations and management functionality.

The following things could be done to extend the scope of this research and provide more complete results.

- Real-world datasets of IPsec rules could be used to both identify potential shortcomings with the approaches presented here and validate their ability to represent the data using more than one dataset.

- Additional visualization approaches could be evaluated to possibly identify more efficient ways of representing IPsec data.

- Additional approaches to the visualizations presented here could be explored to further improve their clarity and the effectiveness of the display.

- A user study involving people of varying backgrounds and experience levels to evaluate various data representations and interface features could help validate the results presented here and identify additional issues or potential functionality.

- A survey of existing IPsec management tools and network management tools could provide insight into additional management functionality or interface options.

- The use of multiple views, particularly pairing the interface proposed here with an existing network management visualization, may allow an administrator to maintain a higher level of situation awareness and provided new directions for effective management functionality not possible otherwise.

5.5. Summary

As a result of the research presented in this thesis, it is possible to develop a tool that could simplify IPsec management using visualization techniques to present IPsec information. Visualization makes it possible to convey large amounts of data quickly and clearly, and it allows a viewer to focus on items of interest while pushing additional data (noise) to the background or hiding it altogether. Additionally, since there are many ways to represent the same data, multiple views of the data can be presented to the viewer allowing for a high degree of adaptability. The visual approaches also present new directions for developing management functionality beyond what is currently available in today's IPsec management tools. If fully realized, this visual approach could simplify IPsec management, making this powerful tool more accessible and more viable for use on production networks. Additionally, the approaches outlined in this thesis could be applied beyond IPsec to any system that similarly defines a system of logical network connections.

Appendix A: IPsec Rules

This appendix provides the consolidated list of IPsec rules applied to each node in the network used to create the visualizations in Chapter IV. When looking at the complete set of rules (included electronically), each rule in the table is repeated four times. For a rule between Host A and Host B, there is an outgoing and incoming rule configured on both hosts. However, since they are duplicates, each rule only needs to be rendered once. The table below is a consolidation of the IPsec rules for all 32 nodes with the duplicate rules removed. This is the table that was used to produce the simulation in Prefuse.

Consolidated IPsec rules

Rule #	Src Add	Dest Add	Src Port	Dest Port	Protocol	Action	Enc Alg	Auth Alg
1	.1.2	.1.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
2	.1.2	.1.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
3	.1.2	.1.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
4	.1.2	.10.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
5	.1.2	.10.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
6	.1.2	.10.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
7	.1.2	.10.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
8	.1.2	.10.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
9	.1.2	.20.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
10	.1.2	.20.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
11	.1.2	.20.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
12	.1.2	.20.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
13	.1.2	.20.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
14	.1.2	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
15	.1.2	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
16	.1.2	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
17	.1.2	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
18	.1.2	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
19	.1.2	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
20	.1.2	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
21	.1.2	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
22	.1.2	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
23	.1.2	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
24	.1.2	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
25	.1.2	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1

26	.1.2	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
27	.1.2	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
28	.1.2	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
29	.1.2	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
30	.1.3	.1.4	53	53	UDP	IPsec	NULL	HMAC-SHA-1
31	.1.3	.1.4	53	53	TCP	IPsec	NULL	HMAC-SHA-1
32	.1.3	.5.2	53	53	UDP	IPsec	NULL	HMAC-SHA-1
33	.1.3	.5.2	53	53	TCP	IPsec	NULL	HMAC-SHA-1
34	.1.3	.10.1	53	53	UDP	IPsec	NULL	HMAC-SHA-1
35	.1.3	.10.2	53	53	UDP	IPsec	NULL	HMAC-SHA-1
36	.1.3	.10.3	53	53	UDP	IPsec	NULL	HMAC-SHA-1
37	.1.3	.10.4	53	53	UDP	IPsec	NULL	HMAC-SHA-1
38	.1.3	.10.5	53	53	UDP	IPsec	NULL	HMAC-SHA-1
39	.1.3	.20.1	53	53	UDP	IPsec	NULL	HMAC-SHA-1
40	.1.3	.20.2	53	53	UDP	IPsec	NULL	HMAC-SHA-1
41	.1.3	.20.3	53	53	UDP	IPsec	NULL	HMAC-SHA-1
42	.1.3	.20.4	53	53	UDP	IPsec	NULL	HMAC-SHA-1
43	.1.3	.20.5	53	53	UDP	IPsec	NULL	HMAC-SHA-1
44	.1.3	.30.1	53	53	UDP	IPsec	NULL	HMAC-SHA-1
45	.1.3	.30.2	53	53	UDP	IPsec	NULL	HMAC-SHA-1
46	.1.3	.30.3	53	53	UDP	IPsec	NULL	HMAC-SHA-1
47	.1.3	.30.4	53	53	UDP	IPsec	NULL	HMAC-SHA-1
48	.1.3	.30.5	53	53	UDP	IPsec	NULL	HMAC-SHA-1
49	.1.3	.30.6	53	53	UDP	IPsec	NULL	HMAC-SHA-1
50	.1.3	.30.7	53	53	UDP	IPsec	NULL	HMAC-SHA-1
51	.1.3	.40.1	53	53	UDP	IPsec	NULL	HMAC-SHA-1
52	.1.3	.40.2	53	53	UDP	IPsec	NULL	HMAC-SHA-1
53	.1.3	.40.3	53	53	UDP	IPsec	NULL	HMAC-SHA-1
54	.1.3	.40.4	53	53	UDP	IPsec	NULL	HMAC-SHA-1
55	.1.3	.40.5	53	53	UDP	IPsec	NULL	HMAC-SHA-1
56	.1.3	.40.6	53	53	UDP	IPsec	NULL	HMAC-SHA-1
57	.1.3	.40.7	53	53	UDP	IPsec	NULL	HMAC-SHA-1
58	.1.3	.40.8	53	53	UDP	IPsec	NULL	HMAC-SHA-1
59	.1.3	.40.9	53	53	UDP	IPsec	NULL	HMAC-SHA-1
60	.1.3	.10.1	53	53	TCP	IPsec	NULL	HMAC-SHA-1
61	.1.3	.10.2	53	53	TCP	IPsec	NULL	HMAC-SHA-1
62	.1.3	.10.3	53	53	TCP	IPsec	NULL	HMAC-SHA-1
63	.1.3	.10.4	53	53	TCP	IPsec	NULL	HMAC-SHA-1
64	.1.3	.10.5	53	53	TCP	IPsec	NULL	HMAC-SHA-1
65	.1.3	.20.1	53	53	TCP	IPsec	NULL	HMAC-SHA-1
66	.1.3	.20.2	53	53	TCP	IPsec	NULL	HMAC-SHA-1
67	.1.3	.20.3	53	53	TCP	IPsec	NULL	HMAC-SHA-1
68	.1.3	.20.4	53	53	TCP	IPsec	NULL	HMAC-SHA-1
69	.1.3	.20.5	53	53	TCP	IPsec	NULL	HMAC-SHA-1
70	.1.3	.30.1	53	53	TCP	IPsec	NULL	HMAC-SHA-1
71	.1.3	.30.2	53	53	TCP	IPsec	NULL	HMAC-SHA-1
72	.1.3	.30.3	53	53	TCP	IPsec	NULL	HMAC-SHA-1

73	.1.3	.30.4	53	53	TCP	IPsec	NULL	HMAC-SHA-1
74	.1.3	.30.5	53	53	TCP	IPsec	NULL	HMAC-SHA-1
75	.1.3	.30.6	53	53	TCP	IPsec	NULL	HMAC-SHA-1
76	.1.3	.30.7	53	53	TCP	IPsec	NULL	HMAC-SHA-1
77	.1.3	.40.1	53	53	TCP	IPsec	NULL	HMAC-SHA-1
78	.1.3	.40.2	53	53	TCP	IPsec	NULL	HMAC-SHA-1
79	.1.3	.40.3	53	53	TCP	IPsec	NULL	HMAC-SHA-1
80	.1.3	.40.4	53	53	TCP	IPsec	NULL	HMAC-SHA-1
81	.1.3	.40.5	53	53	TCP	IPsec	NULL	HMAC-SHA-1
82	.1.3	.40.6	53	53	TCP	IPsec	NULL	HMAC-SHA-1
83	.1.3	.40.7	53	53	TCP	IPsec	NULL	HMAC-SHA-1
84	.1.3	.40.8	53	53	TCP	IPsec	NULL	HMAC-SHA-1
85	.1.3	.40.9	53	53	TCP	IPsec	NULL	HMAC-SHA-1
86	.1.4	.10.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
87	.1.4	.10.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
88	.1.4	.10.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
89	.1.4	.10.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
90	.1.4	.10.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
91	.1.4	.20.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
92	.1.4	.20.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
93	.1.4	.20.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
94	.1.4	.20.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
95	.1.4	.20.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
96	.1.4	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
97	.1.4	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
98	.1.4	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
99	.1.4	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
100	.1.4	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
101	.1.4	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
102	.1.4	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
103	.1.4	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
104	.1.4	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
105	.1.4	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
106	.1.4	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
107	.1.4	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
108	.1.4	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
109	.1.4	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
110	.1.4	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
111	.1.4	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
112	.1.5	.10.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
113	.1.5	.10.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
114	.1.5	.10.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
115	.1.5	.10.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
116	.1.5	.10.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
117	.1.5	.20.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
118	.1.5	.20.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
119	.1.5	.20.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1

120	.1.5	.20.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
121	.1.5	.20.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
122	.1.5	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
123	.1.5	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
124	.1.5	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
125	.1.5	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
126	.1.5	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
127	.1.5	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
128	.1.5	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
129	.1.5	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
130	.1.5	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
131	.1.5	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
132	.1.5	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
133	.1.5	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
134	.1.5	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
135	.1.5	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
136	.1.5	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
137	.1.5	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
138	.5.2	.5.3	53	53	UDP	IPSec	NULL	HMAC-SHA-1
139	.5.2	.5.3	53	53	TCP	IPSec	NULL	HMAC-SHA-1
140	.5.3	.10.1	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
141	.5.3	.10.2	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
142	.5.3	.10.3	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
143	.5.3	.10.4	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
144	.5.3	.10.5	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
145	.5.3	.20.1	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
146	.5.3	.20.2	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
147	.5.3	.20.3	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
148	.5.3	.20.4	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
149	.5.3	.20.5	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
150	.5.3	.30.1	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
151	.5.3	.30.2	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
152	.5.3	.30.3	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
153	.5.3	.30.4	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
154	.5.3	.30.5	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
155	.5.3	.30.6	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
156	.5.3	.30.7	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
157	.5.3	.40.1	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
158	.5.3	.40.2	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
159	.5.3	.40.3	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
160	.5.3	.40.4	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
161	.5.3	.40.5	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
162	.5.3	.40.6	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
163	.5.3	.40.7	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
164	.5.3	.40.8	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
165	.5.3	.40.9	80	ANY	HTTP	IPsec	NULL	HMAC-SHA-1
166	.5.3	.10.1	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1

167	.5.3	.10.2	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
168	.5.3	.10.3	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
169	.5.3	.10.4	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
170	.5.3	.10.5	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
171	.5.3	.20.1	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
172	.5.3	.20.2	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
173	.5.3	.20.3	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
174	.5.3	.20.4	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
175	.5.3	.20.5	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
176	.5.3	.30.1	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
177	.5.3	.30.2	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
178	.5.3	.30.3	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
179	.5.3	.30.4	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
180	.5.3	.30.5	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
181	.5.3	.30.6	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
182	.5.3	.30.7	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
183	.5.3	.40.1	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
184	.5.3	.40.2	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
185	.5.3	.40.3	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
186	.5.3	.40.4	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
187	.5.3	.40.5	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
188	.5.3	.40.6	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
189	.5.3	.40.7	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
190	.5.3	.40.8	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
191	.5.3	.40.9	443	ANY	HTTPS	IPsec	TripleDES-CBC	HMAC-SHA-1
192	.10.1	.20.1	5566	5566	TCP	IPsec	TripleDES-CBC	HMAC-SHA-1
193	.10.1	.30.1	5566	5566	TCP	IPsec	TripleDES-CBC	HMAC-SHA-1
194	.10.1	.40.1	5566	5566	TCP	IPsec	TripleDES-CBC	HMAC-SHA-1
195	.10.1	.10.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
196	.10.1	.10.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
197	.10.1	.10.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
198	.10.1	.10.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
199	.10.1	.20.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
200	.10.1	.20.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
201	.10.1	.20.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
202	.10.1	.20.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
203	.10.1	.20.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
204	.10.1	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
205	.10.1	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
206	.10.1	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
207	.10.1	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
208	.10.1	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
209	.10.1	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
210	.10.1	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
211	.10.1	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
212	.10.1	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
213	.10.1	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1

214	.10.1	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
215	.10.1	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
216	.10.1	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
217	.10.1	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
218	.10.1	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
219	.10.1	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
220	.10.2	.40.2	9000	9000	TCP	IPsec	TripleDES-CBC	HMAC-SHA-1
221	.10.2	.40.3	9000	9000	TCP	IPsec	TripleDES-CBC	HMAC-SHA-1
222	.10.2	.40.4	9000	9000	TCP	IPsec	TripleDES-CBC	HMAC-SHA-1
223	.10.2	.40.5	9000	9000	TCP	IPsec	TripleDES-CBC	HMAC-SHA-1
224	.10.2	.40.6	9000	9000	TCP	IPsec	TripleDES-CBC	HMAC-SHA-1
225	.10.2	.10.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
226	.10.2	.10.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
227	.10.2	.10.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
228	.10.2	.20.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
229	.10.2	.20.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
230	.10.2	.20.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
231	.10.2	.20.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
232	.10.2	.20.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
233	.10.2	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
234	.10.2	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
235	.10.2	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
236	.10.2	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
237	.10.2	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
238	.10.2	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
239	.10.2	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
240	.10.2	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
241	.10.2	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
242	.10.2	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
243	.10.2	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
244	.10.2	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
245	.10.2	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
246	.10.2	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
247	.10.2	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
248	.10.2	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
249	.10.3	.20.2	ANY	ANY	ANY	IPsec	TripleDES-CBC	HMAC-SHA-1
250	.10.3	.20.3	ANY	ANY	ANY	IPsec	TripleDES-CBC	HMAC-SHA-1
251	.10.3	.10.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
252	.10.3	.10.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
253	.10.3	.20.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
254	.10.3	.20.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
255	.10.3	.20.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
256	.10.3	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
257	.10.3	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
258	.10.3	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
259	.10.3	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
260	.10.3	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1

261	.10.3	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
262	.10.3	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
263	.10.3	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
264	.10.3	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
265	.10.3	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
266	.10.3	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
267	.10.3	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
268	.10.3	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
269	.10.3	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
270	.10.3	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
271	.10.3	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
272	.10.4	.10.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
273	.10.4	.20.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
274	.10.4	.20.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
275	.10.4	.20.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
276	.10.4	.20.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
277	.10.4	.20.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
278	.10.4	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
279	.10.4	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
280	.10.4	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
281	.10.4	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
282	.10.4	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
283	.10.4	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
284	.10.4	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
285	.10.4	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
286	.10.4	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
287	.10.4	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
288	.10.4	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
289	.10.4	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
290	.10.4	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
291	.10.4	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
292	.10.4	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
293	.10.4	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
294	.10.5	.20.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
295	.10.5	.20.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
296	.10.5	.20.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
297	.10.5	.20.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
298	.10.5	.20.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
299	.10.5	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
300	.10.5	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
301	.10.5	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
302	.10.5	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
303	.10.5	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
304	.10.5	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
305	.10.5	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
306	.10.5	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
307	.10.5	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1

308	.10.5	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
309	.10.5	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
310	.10.5	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
311	.10.5	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
312	.10.5	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
313	.10.5	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
314	.10.5	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
315	.20.1	.30.1	5566	5566	TCP	IPsec	TripleDES-CBC	HMAC-SHA-1
316	.20.1	.40.1	5566	5566	TCP	IPsec	TripleDES-CBC	HMAC-SHA-1
317	.20.1	.20.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
318	.20.1	.20.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
319	.20.1	.20.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
320	.20.1	.20.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
321	.20.1	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
322	.20.1	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
323	.20.1	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
324	.20.1	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
325	.20.1	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
326	.20.1	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
327	.20.1	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
328	.20.1	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
329	.20.1	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
330	.20.1	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
331	.20.1	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
332	.20.1	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
333	.20.1	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
334	.20.1	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
335	.20.1	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
336	.20.1	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
337	.20.2	.10.3	ANY	ANY	ANY	IPsec	TripleDES-CBC	HMAC-SHA-1
338	.20.2	.20.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
339	.20.2	.20.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
340	.20.2	.20.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
341	.20.2	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
342	.20.2	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
343	.20.2	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
344	.20.2	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
345	.20.2	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
346	.20.2	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
347	.20.2	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
348	.20.2	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
349	.20.2	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
350	.20.2	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
351	.20.2	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
352	.20.2	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
353	.20.2	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
354	.20.2	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1

355	.20.2	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
356	.20.2	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
357	.20.3	.10.3	ANY	ANY	ANY	IPsec	TripleDES-CBC	HMAC-SHA-1
358	.20.3	.20.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
359	.20.3	.20.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
360	.20.3	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
361	.20.3	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
362	.20.3	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
363	.20.3	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
364	.20.3	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
365	.20.3	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
366	.20.3	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
367	.20.3	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
368	.20.3	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
369	.20.3	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
370	.20.3	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
371	.20.3	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
372	.20.3	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
373	.20.3	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
374	.20.3	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
375	.20.3	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
376	.20.4	.20.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
377	.20.4	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
378	.20.4	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
379	.20.4	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
380	.20.4	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
381	.20.4	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
382	.20.4	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
383	.20.4	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
384	.20.4	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
385	.20.4	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
386	.20.4	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
387	.20.4	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
388	.20.4	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
389	.20.4	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
390	.20.4	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
391	.20.4	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
392	.20.4	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
393	.20.5	.30.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
394	.20.5	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
395	.20.5	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
396	.20.5	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
397	.20.5	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
398	.20.5	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
399	.20.5	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
400	.20.5	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
401	.20.5	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1

402	.20.5	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
403	.20.5	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
404	.20.5	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
405	.20.5	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
406	.20.5	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
407	.20.5	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
408	.20.5	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
409	.30.1	.40.1	5566	5566	TCP	IPsec	TripleDES-CBC	HMAC-SHA-1
410	.30.1	.30.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
411	.30.1	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
412	.30.1	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
413	.30.1	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
414	.30.1	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
415	.30.1	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
416	.30.1	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
417	.30.1	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
418	.30.1	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
419	.30.1	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
420	.30.1	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
421	.30.1	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
422	.30.1	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
423	.30.1	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
424	.30.1	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
425	.30.2	.30.3	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
426	.30.2	.30.4	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
427	.30.2	.30.5	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
428	.30.2	.30.6	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
429	.30.2	.30.7	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
430	.30.2	.30.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
431	.30.2	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
432	.30.2	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
433	.30.2	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
434	.30.2	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
435	.30.2	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
436	.30.2	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
437	.30.2	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
438	.30.2	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
439	.30.2	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
440	.30.2	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
441	.30.2	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
442	.30.2	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
443	.30.2	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
444	.30.3	.30.4	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
445	.30.3	.30.5	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
446	.30.3	.30.6	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
447	.30.3	.30.7	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
448	.30.3	.30.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1

449	.30.3	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
450	.30.3	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
451	.30.3	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
452	.30.3	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
453	.30.3	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
454	.30.3	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
455	.30.3	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
456	.30.3	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
457	.30.3	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
458	.30.3	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
459	.30.3	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
460	.30.3	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
461	.30.4	.30.5	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
462	.30.4	.30.6	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
463	.30.4	.30.7	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
464	.30.4	.30.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
465	.30.4	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
466	.30.4	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
467	.30.4	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
468	.30.4	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
469	.30.4	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
470	.30.4	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
471	.30.4	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
472	.30.4	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
473	.30.4	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
474	.30.4	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
475	.30.4	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
476	.30.5	.30.6	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
477	.30.5	.30.7	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
478	.30.5	.30.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
479	.30.5	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
480	.30.5	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
481	.30.5	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
482	.30.5	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
483	.30.5	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
484	.30.5	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
485	.30.5	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
486	.30.5	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
487	.30.5	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
488	.30.5	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
489	.30.6	.30.7	23	23	TELNET	IPsec	TripleDES-CBC	HMAC-SHA-1
490	.30.6	.30.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
491	.30.6	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
492	.30.6	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
493	.30.6	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
494	.30.6	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
495	.30.6	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1

496	.30.6	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
497	.30.6	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
498	.30.6	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
499	.30.6	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
500	.30.7	.40.1	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
501	.30.7	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
502	.30.7	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
503	.30.7	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
504	.30.7	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
505	.30.7	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
506	.30.7	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
507	.30.7	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
508	.30.7	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
509	.40.1	.40.2	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
510	.40.1	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
511	.40.1	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
512	.40.1	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
513	.40.1	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
514	.40.1	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
515	.40.1	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
516	.40.1	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
517	.40.2	.40.3	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
518	.40.2	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
519	.40.2	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
520	.40.2	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
521	.40.2	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
522	.40.2	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
523	.40.2	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
524	.40.3	.40.4	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
525	.40.3	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
526	.40.3	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
527	.40.3	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
528	.40.3	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
529	.40.3	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
530	.40.4	.40.5	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
531	.40.4	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
532	.40.4	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
533	.40.4	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
534	.40.4	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
535	.40.5	.40.6	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
536	.40.5	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
537	.40.5	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
538	.40.5	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
539	.40.6	.40.7	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
540	.40.6	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
541	.40.6	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
542	.40.7	.40.8	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1

543	.40.7	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1
542	.40.8	.40.9	ANY	ANY	ANY	IPsec	NULL	HMAC-SHA-1

Bibliography

[1] "Request for Comments 4301: Security Architecture for the Internet Protocol", Internet Engineering Task Force, November, 2005. URL http://www.ietf.org/rfc/rfc4301.txt.

[2] Ware, Colin, *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers, San Francisco, California, 2004. ISBN 1-55860-819-2.

[3] URL http://www-personal.umich.edu/~mejn/election/2008/.

[4] Conti, Greg, *Security Data Visualization: Graphical Techniques for Network Analysis*. No Starch Press, San Francisco, California, 2007. ISBN 1-59327-143-3.

[5] Marty, Raffael, *Applied Security Visualization*. Addison-Wesley Professional, 2008. ISBN 0-321-51010-0.

[6] URL http://www.infovis.net/printMag.php?lang=2&num=179.

[7] URL http://www.cs.umd.edu/class/spring2005/cmsc838s/viz4all/viz4all_a.html.

[8] URL http://www.logixml.com/images/products/pops/info/visualization.jpg.

[9] Tufte, Edward, *The Visual Display of Quantitative Information, 2nd ed*. Graphics Press LLC, Cheshire, Connecticut, 2001. ISMB 978-0-9613921-4-7.

[10] Few, Stephen, "An Introduction to Visual Multivariate Analysis", Perceptual Edge, July 11, 2006. URL http://www.perceptualedge.com/articles/b-eye/visual_multivariate_analysis.pdf.

[11] Few, Stephen, "Multivariate Analysis Using Parallel Coordinates", Perceptual Edge, September 12, 2006. URL http://www.perceptualedge.com/articles/b-eye/parallel_coordinates.pdf.

[12] URL http://farm2.static.flickr.com/1343/775147043_b28938bb3e_o.gif.

[13] Card, S., Mackinlay, J., Shneiderman, B., *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers, San Francisco, California, 1999. ISBN 1-55860-533-9.

[14] URL http://www.cs.umd.edu/hcil/treemap-history/.

[15] Endsley, Mica, "Toward a Theory of Situation Awareness in Dynamic Systems", in *Human Factors Journal*, Volume 37(1), pages 32-64, March 1995.

[16] Endsley, Mica, "Theoretical Underpinnings of Situation Awareness: A Critical Review", in *Situation Awareness Analysis and Measurement*, pages 3-32, Lawrence Erlbaum Associates Inc.

[17] Tadda, G., Boulware, D., Hinman, M., Gorton, S, "Realizing Situation Awareness in a Cyber Environment", in *Multisensor, Multisource Information Fusion: Architechtures, Algorithms, and Applications* 2006, Proceedings of SPIE Vol. 6242 (SPIE, Bellingham, WA, 2006) 624204.

[18] Okolica, J., McDonald, T., Peterson, G., Mills, R., Haas, M., "Developing Systems for Cyber Situational Awareness", In Proceedings of the 2nd Cyberspace Research Workshop, June, 2009..

[19] Bibighaus, David, "Cybercraft Whitepaper", Air Force Research Lab, Rome, NY, September 2006.

[20] Kozierok, Charles, *The TCP/IP Guide: A Comprehensive, Illustrated Internet Protocols Reference*. No Starch Press, 2005. ISBN 159327047X.

[21] "Request for Comments 4302: IP Authentication Header", Internet Engineering Task Force, November, 2005. URL http://www.ietf.org/rfc/rfc4302.txt.

[22] "Request for Comments 4303: IP Encapsulating Security Payload", Internet Engineering Task Force, November, 2005. URL http://www.ietf.org/rfc/rfc4303.txt.

[23] "Request for Comments 4306: Internet Key Exchange (IKEv2) Protocol", Internet Engineering Task Force, November, 2005. URL http://www.ietf.org/rfc/rfc4306.txt.

[24] Frankel, Sheila, *Demystifying the IPsec Puzzle.* Artech House, Inc., Norwood, Massachusetts, 2001. ISBN 1-58053-079-6.

[25] Doraswamy, Naganand and Harkins, Dan, *IPSec: The New Security Standard for the Internet, Intranets, and Virtual Private Networks*. Prentice Hall PTR, Upper Saddle River, New Jersey, 2003. ISBN 0-13-046189.

[26] Snader, Jon C., *VPNs Illustrated: Tunnels, VPNs, and IPsec*. Addison-Wesley Professional, 2005. ISBN 0-321-24544-X.

[27] "Request for Comments 4305: Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authenticiation Header (AH)", Internet Engineering Task Force, December, 2005. URL http://www.ietf.org/rfc/rfc4305.txt.

[28] NIST, Secure Hash Standard. FIPS Pub. 180-2, April 1995. URL http://csrc.nist.gov/publications/fips/fips180-2/fips180-2withchangenotice.pdf.

[29] "Request for Comments 2451: The ESP CBC-Mode Cipher Algorithms", Internet Engineering Task Force, November, 1998. URL http://www.ietf.org/rfc/rfc2451.txt.

[30] "Request for Comments 3602: The AES-CBC Cipher Algorithm and Its Use with IPsec", Internet Engineering Task Force, September, 2003. URL http://www.ietf.org/rfc/rfc3602.txt.

[31] URL http://docs.hp.com/en/J4256-90015/ch01s02.html.

[32] Ferguson, Niels and Schneier, Bruce, "A Cryptographic Evaluation of IPsec", February, 1999. URL http://www.counterpane.com/ipsec.html.

[33] Tiller, James, *A Technical Guide to IPSec Virtual Private Networks*, Auerbach, December, 2000. ISBN 0849308763.

[34] URL http://prefuse.org/.

[35] URL http://www.solarwinds.com/products/LANsurveyor.

[36] Conti, Greg. "Thesis Correspondence." Electronic Message. 26 October 2009.

[37] "Request for Comments 2408: Internet Security Association and Key Management Protocol (ISAKMP) ", Internet Engineering Task Force, November, 1998. URL http://tools.ietf.org/html/rfc2408.

[38] "Request for Comments 2412: The OAKLEY Key Determination Protocol ", Internet Engineering Task Force, November, 1998. URL http://tools.ietf.org/html/rfc2412.

[39] Krawczyk, H., "SKEME: A Versatile Secure Key Exchange Mechanism for Internet", from IEEE Proceedings of the 1996 Symposium on Network and Distributed Systems Security. URL www.isoc.org/isoc/conferences/ndss/96/ndss96/krawczyk_slides.ps.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From – To)*
08-02-2010	**Master's Thesis**	09-2008 – 03-2010

4. TITLE AND SUBTITLE

Visually Managing IPsec

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Dell'Accio, Peter J., 1st Lt, USAF

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S)

Air Force Institute of Technology
Graduate School of Engineering and Management (AFIT/EN)
2950 Hobson Way
WPAFB OH 45433-7765

8. PERFORMING ORGANIZATION REPORT NUMBER

AFIT/GCO/ENG/10-06

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

INTENTIONALLY LEFT BLANK

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The United States Air Force relies heavily on computer networks to transmit vast amounts of information throughout its organizations and with agencies throughout the Department of Defense. The data take many forms, utilize different protocols, and originate from various platforms and applications. It is not practical to apply security measures specific to individual applications, platforms, and protocols. Internet Protocol Security (IPsec) is a set of protocols designed to secure data traveling over IP networks, including the Internet. By applying security at the network layer of communications, data packets can be secured regardless of what application generated the data or which protocol is used to transport it. However, the complexity of managing IPsec on a production network, particularly using the basic command-line tools available today, is the limiting factor to widespread deployment. This thesis explores several visualizations of IPsec data, evaluates the viability of using visualization to represent and manage IPsec, and proposes an interface for a visual IPsec management application to simplify IPsec management and make this powerful security option more accessible to the information warfighter.

15. SUBJECT TERMS

Visualization, IPsec, Network Management, Network Security

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
REPORT	**ABSTRACT**	**c. THIS PAGE**			**Lt Col Stuart H. Kurkowski (ENG)**
U	**U**	**U**	**UU**	146	**19b. TELEPHONE NUMBER** *(Include area code)* (937) 255-3636 stuart.kurkowski@afit.edu

Standard Form 298 (Rev: 8-98)
Prescribed by ANSI Std. Z39-18